Comparative Review: Server Defrag p. 25 Buyer's Guide: Enterprise Antivirus p. 29

MANAGEMENT • SECURITY • NETWORKING • MESSAGING • HOW-TO

Windows IT Pro

WE'RE IN IT

May 29, 2008
EXCHANGE 2007
Mastery Series
Server Management
PAGE 55

Master RSoP: MAP GROUP POLICY MUTATIONS

p. 32

Manage PST Files in Outlook p. 43

Integrate AD and OpenLDAP p. 46

PowerShell: Get Your Quotes Right p. 51

Microsoft's Brad Anderson

Microsoft's Larry Orecklin

EXCLUSIVE! Groundbreaking System Center News p. 39

OFFICE AND SHAREPOINT PRO

Improve Collaboration with Exchange and MOSS p. 57

Control SharePoint Access p. 61

Penton
A PENTON PUBLICATION
MAY 2008
WWW.WINDOWSITPRO.COM
U.S. $5.95 CANADA $7.95

Windows IT Pro
MAY 2008 VOLUME 14 NO. 5

FEATURES

SOLUTIONS+

COVER STORY

IT PRO HERO

FEATURES

OFFICE & SHAREPOINT PRO

TRICKS & TRAPS

COLUMNS

Karen Forster
IT Pro Perspective

Windows Vista and Windows Server 2008 Togetherness

Will the "better together" features and common code base motivate you to speed up your migration to the new client and server OSs? Microsoft's product managers for Windows Server and Windows Client Deployment discuss how upgrading to Server 2008 and Vista together improve performance, security, and manageability.

Paul Thurrott
Need to Know

What You Need to Know About Microsoft Response Point and Windows Vista SP1 Scheduling

Microsoft's small business phone system, Response Point, is worth a try—especially after the SP1 update offering VoIP integration. Discover the reason for Windows Vista SP1's strange release schedule.

Mark Minasi
Windows Power Tools

Control Windows Server 2008 Roles and Features

You can work around Server 2008's lack of a Computer Management console. Efficiently configure your servers and roll them out by using the ServerManagerCmd, Ocsetup, and Oclist command-line tools.

Michael Otey
Top 10

Vista Command-Line Tools

Modify your boot configuration, manage ACLs, perform backups and system assessments, and take advantage of other useful functions with these helpful Windows Vista command-line tools.

PRODUCTS

WHAT'S HOT

Jeff Kowalke, senior Microsoft systems engineer

IN EVERY ISSUE

windowsitpro.com

Making the Case for Oracle Database on Windows

Find out if running Oracle on Windows is the right combination for your organization by walking through the migration process with Kroll Factual Data. Read about the challenges Kroll Factual Data faced, why the company decided to migrate to Oracle, and the core integration capabilities that Oracle offers to users deploying its databases and using its tools within Windows-based environments.

www.windowsitpro.com/go/wp/ovum/oracle/?code=maycitc

YOUR SAVVY ASSISTANT

The Missing Link to IT Resources

Microsoft Gets Personal

Microsoft wants to know who you are (i.e., wants you to do silly things and then post a video of them on its "Who Are You?" Web site—www.wewantoknow.net). Karen Forster explains the company's campaign in her December 2007 article "Microsoft Asks: Who Are You?" (InstantDoc ID 97478). She refers to a flyer that explains the program as a place "where IT Professionals can showcase their multidimensional, creative personalities as people instead of simply professionals." I understand the importance of showing people as more than just their job titles, but this attempt misses the mark.

Pardon me for calling 'em like I see 'em, but other than the guy with the super-smart kid and the montage of people being knocked in the nether regions, the Who Are You? videos I've watched show only one dimension: karaoke. I love off-key renditions of Neil Diamond tunes just as much as the next guy, but I just don't see how they're useful for your career.

Reader Curt Hayes seems to agree with me. In a February 2008 letter to the editor (InstantDoc ID 97859), Curt writes, "Maybe I'm just old and grouchy, but I don't see how celebrating my ability to play the kazoo translates into helping me do my job." Although I would love to hear Curt's kazoo talents, I'd rather help him get through his workday more efficiently so that he could spend more time doing what he enjoys. And so should Microsoft.

If you want to connect with peers as a person and a professional, check out our forums at forums.windowsitpro.com. And if you participate in Facebook or MySpace, check out the *Windows IT Pro* profiles at www.windowsitpro.com/go/facebook and www.myspace.com/windowsitpro and Your Savvy Assistant's profile at www.myspace.com/yoursavvyassistant. Of course, I always welcome your questions, feedback, and photos (e.g., reader Tim Harper—pictured here with a very cheesy caption) at christan.humphries@penton.com.

Water You Waiting for? Show Us Who You Really Are!

Step-by-Step Guide to Disaster Recovery Planning

Combine available backup and recovery technologies to find a holistic approach to developing, implementing, and testing your disaster recovery plan. Outline the steps you should follow to ensure that your disaster recovery plan will work as you expect it to and that it will scale as your organization and IT needs evolve.

www.windowsitpro.com/go/seminars/XOsoft/DisasterRecovery/?partnerref=maycitc

Creating Data Storage Systems that Meet Regulatory Guidelines

Ensure that your data storage system meets regulatory guidelines. Regulatory compliance is often a major bugaboo for the storage administrator. In this podcast, David Chernicoff talks about the regulatory issues surrounding data storage requirements, such as the Sarbanes-Oxley (SOX) Act of 2002, SEC Rule 17-4, and The Health Insurance Portability and Accountability Act (HIPAA). He will also provide direction to IT personnel who are responsible for setting up storage systems that meet regulatory guidelines.

www.windowsitpro.com/go/storageguardian/pod/regulatory/?code=maycitc

Find more information about disaster recovery plans and storage systems on the Storage/Backup/Recovery forum.

www.windowsitpro.com/go/forum/StorageBackupRecovery

It's Not Just E-mail.

Lucid
www.Lucid8.com

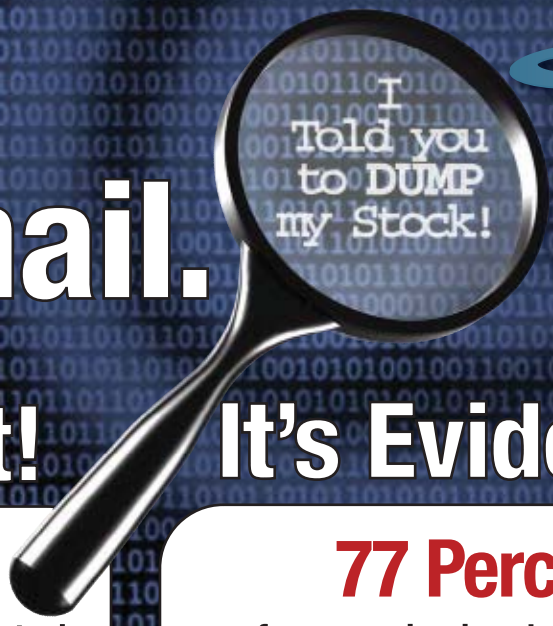

It's an Asset! It's Evidence!

75 Percent
of corporate intellectual property is sent through email messages and their attachments.*

77 Percent
of companies involved in legal or regulatory actions had email requested as part of the discovery process.*

• DISCOVER • RECOVER • EXPORT

DISCOVER: Create and reuse advanced queries to search a single data source or across multiple **Backup Copies of Exchange Information Stores or Live Exchange Servers** as well as PST's and DigiVault data sets to find the required evidence within emails, attachments and meta-data.

RECOVER: Use DigiScope's intuitive Outlook interface to restore information via drag-&-drop to a specific location or select SingleTouch™ recovery to automatically restore mailboxes, folders, or individual items to original locations within the live Exchange Server.

EXPORT: Search results can be optionally de-duplicated and then exported to multiple formats including, XML, MSG, and PST's with various options to support data migration as well as further review or legal analysis.

LIVE WEEKLY DEMOS
powered by webex

FREE DOWNLOADS
- Demo version of DigiScope
- White Papers
  - FRCP, E-mail Discovery & You
  - Essential Guide to E-discovery and Recovery

Go to:
www.Lucid8.com/WIPdiscover
Call: **425 456-8496**
E-Mail: **Sales@Lucid8.com**

Lucid8's **DigiScope**

eDiscovery and Recovery for Microsoft® Exchange

*Source: Enterprise Strategy Group

NINJA™ BLADE

The Barracuda Killer.

Carrier-grade Email Security on Dell Hardware

Introducing Ninja Blade. Want the most affordable, best-in-class, carrier-grade email security appliances on the market today? Then Ninja Blade is your choice. It is built on commercial software, as opposed to inexpensive solutions which re-brand open source offerings. The hardware is industrial-grade as well, built through a partnership with Dell Inc., with one year of onsite service included with every purchase. What you get is a powerful, industrial-strength appliance, at a price point that beats the other guys.

Best-of-breed antispam and antivirus protection. For antispam protection, Ninja Blade uses the Cloudmark™ antispam engine which provides superior signature-based detection against spam, phishing, and email spoofing. To protect against viruses and zero-day threats, Ninja Blade incorporates best-in-class antivirus technology from BitDefender™. BitDefender's antivirus engine provides comprehensive protection against viruses, backdoors, trojans, and other malicious threats, certified by Virus Bulletin, ICSA Labs and Checkmark.

Uses the MessageSystems carrier-class MTA. The foundation of Ninja Blade is its use of MessageSystems' carrier-class MTA (Message Transfer Agent), capable of handling millions of unique messages per hour – making it one of the fastest solutions on the market today.

"Ninja Blade is built on commercial software, as opposed to inexpensive solutions which re-brand open source offerings."

Comprehensive, rules-based email management. Ninja Blade's management tools allow administrators to easily implement rules for both inbound and outbound email. A variety of rules can be created based on a number of email message properties such as body, sender IP, header, or subject, with the ability to add further custom rules for even more powerful filtering. Ninja Blade also has global attachment filtering and built-in disclaimers.

Easy deployment and administration; integrates with Active Directory. Ninja Blade's easy-to-use, advanced web-based interface enables you to install, configure and manage your entire gateway email security from any web browser. With easy setup and fast deployment, it can be installed and configured in minutes. With real-time Active Directory and LDAP integration, Ninja Blade provides hands-off maintenance and seamlessly integrates with existing IT infrastructures.

Ninja Blade Series for organizations of all sizes. Sunbelt Software is dedicated to providing industry-strength email security products for organizations of all sizes. Ninja Blade 500, 1000, 2500, and 5000 appliances are easy-to-install, and are suitable for organizations with up to 5000 users and multiple domains.

Sunbelt Software

See how Ninja Blade fits into your network today. Contact us for an evaluation unit!

Visit www.ninjablade.com

Sunbelt Software Tel: 1-888-688-8457 or 1-727-562-0101 Fax: 1-727-562-5199 www.sunbeltsoftware.com sales@sunbeltsoftware.com

Windows Vista and Windows Server 2008 Togetherness

Performance, security, and manageability

"The three areas that customers will get the most benefit from by using Windows Server 2008 and Windows Vista together are system-wide performance, improved security, and manageability." Following the launch of Server 2008 and Vista SP1, these were the points emphasized by Microsoft's Justin Graham, senior technical product manager for Windows Server, and his Vista colleague, Jeremy Chapman, senior product manager, Windows Client Deployment. "Within each of these three buckets," Justin continued, "we have a number of key features that the server either enables or makes better when used with the client."

Performance

Because the client and server now share a code base for all versions, from consumer through server, networking is unified and therefore faster. Justin pointed out, "On the networking side, our improvements are on things like the next generation TCP/IP stack. It's completely redesigned. The benefits are native IPv6 support versus emulated IPv6 support in previous versions. Also, some other features that allow things like policy-based Quality of Service [QoS] and Receive Window Auto-Tuning. We also made improvements and have released a new SMB (Server Message Block) protocol called SMB 2.0. This is aimed at improving file sharing performance and also streaming video. The SMB protocol makes all of that happen much, much faster than in previous versions of Windows, including XP SP2."

I asked Justin to elaborate on Receive Window Auto-Tuning. "This feature really benefits companies who have a lot of branch offices or remote offices that have varying speeds of WAN links," Justin replied. "When the client and server are speaking over the network, they can sense the network conditions and automatically increase or decrease their Receive Window to match the conditions of the network—basically enabling the maximum consistent throughput on the network for the given conditions."

Security and Manageability

The features that Microsoft touts as security and manageability advantages to using Vista and Server 2008 together are Network Access Protection (NAP)—which lets you monitor, isolate, and remediate the security of devices as they try to access your network—and Windows BitLocker Drive Encryption. Justin explained, "When you use NAP with Windows Vista, you get a couple of benefits. One is that the NAP client, which ensures that the client is healthy, is automatically included with Windows Vista. With Windows XP, you have to download and deploy an add-on client. That makes it a little more complex to deploy NAP in a Windows XP scenario. The last piece is that with Windows Vista, we support an additional enforcement mechanism called AuthIP. In XP you're limited to deploying IPsec or using hardware-based 802.1x enforcement. The benefit to AuthIP is that it is an extension of IPsec that is more modular and easier to configure." (For information about AuthIP, see technet.microsoft.com/en-us/library/bb878097.aspx.)

As for BitLocker, Justin said, "The real improvement here is we're able to use Group Policy to force BitLocker Drive Encryption on certain groups of client and server machines. With just Windows Vista and not running Windows Server 2008 on the back end, there's no way to force BitLocker Drive Encryption on the clients."

Deployment

Moving the discussion to the client side, Jeremy emphasized "end-to-end deployment," including Windows Deployment Services (WDS) Multicast and Volume Activation. With availability of the Windows Automated Installation Kit (WAIK), Vista and Server 2008 let you take advantage of deployment tools that were previously available only to OEMs to support mass deployments. In addition, Jeremy said, the former Business Desktop Deployment (BDD), now called Microsoft Deployment Toolkit (MDT), provides "task sequencing functionality with very low reliance on scripting. And it will call the new Windows Server 2008 Server Manager, which allows you to apply a role to a server. After that role is applied, we can continue to configure that server to where it's completely usable."

WDS Multicast was a much-requested feature. Jeremy explained, "In the Windows Server 2003 WDS or [Remote Installation Services] RIS days, we were doing one image per device. If you're deploying 1,000 machines and each image is 4GB to 5GB, you're deploying 5,000GB over the network. Instead of 5,000GB, you're doing 10GB. If the machine joins the multicast transmission later, it will even pick up the components it missed during the first transmission. So it might only take two loops to provision that image to 1,000 machines."

Better Together

Microsoft has gone to great lengths to demonstrate the value of implementing Vista and Server 2008 together. Do such features make you more likely to install both new OSs? I'm looking forward to your feedback.

InstantDoc ID 98575

Karen Forster (karen@windowsitpro.com) is group editorial and strategy director for *Windows IT Pro* and *SQL Server Magazine* and former director of Windows Server User Assistance at Microsoft.


Microsoft
In a consolidated IT world, you need servers that run on legs of steel. So we gave Windows Server® 2008 innovations, such as Failover Clustering and a Server Core installation option, that help isolate, resolve, and evade problems to deliver superhuman reliability.
Meet the new Windows Server 2008 at serverunleashed.com
IT 24-7
It understands two things:
"go" and "further."
Windows Server® 2008

busi·ness
pro·cess
au·to·ma·tion

[biz-nis | pros-es | aw-tuh-mey-shuhn]

The replacement of a manual business process with an automated one, usually through the use of **advanced technologies**.

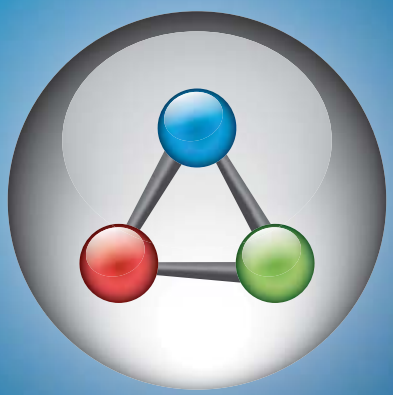

AutoMate BPA Server 7™

The Business Process Automation Server from Network Automation

NO CODE, NO LIMITS

Automates business & IT processes
Eliminates the need for job schedulers, scripts & batch files
Intuitive drag-and-drop workflow design & task development

Visit **WhatIsBPAServer.com** to learn more about **BPA Server 7** and how the world leader in **Business Process Automation** is advancing the field. Again.

NETWORK AUTOMATION

www.WhatIsBPAServer.com
888-786-4796

Redmond MVP Most Valuable Product

Windows IT Pro READERS CHOICE 2006 Winner

Windows IT Pro EDITOR'S BEST 2007 WINNER

Unlocking USB Drives with Cipher

I appreciate Mark Minasi's Cipher tip in his Windows Power Tools column ("Manage Your EFS Keys with Cipher," February 2008, InstantDoc ID 97735). I have a USB drive that I move between two machines. Because of company policy, the contents of any USB drive connected to a laptop are encrypted with Cipher. The encryption occurs during the logoff sequence (a GPO logoff script). The first time, the process can hang the logoff sequense as it processes your USB device (in my case, a 200GB drive at 25 percent capacity). After that, it scans for deltas and runs quicker. Mark's technique let me export/import my keys to the second machine and continue to access the data.

—ajrinaldi

PowerShell Intro

Robert Sheldon's article in the February 2008 issue of Windows IT Pro ("PowerShell 101, Lesson 1," InstantDoc ID 97742) provides a terrific introduction to PowerShell. I would've liked to see some mention of Get-Command and Get-Member because these tools, together with Get-Help, comprise a great selection of tools for learning the capability of PowerShell cmdlets. The article also should have mentioned the fact that PowerShell, as initially installed, can't run scripts. If readers can't get simple scripts to run, they'll conclude that PowerShell doesn't work. The use of Set-ExecutionPolicy, as detailed in the Help file, will overcome this restriction. Robert does give a good overview of aliases and how you can use them. I hope he includes, in follow-up articles, some guidelines about not using them in scripts.

—Richard Siddaway

Thanks for your valuable feedback! You raise a good point about running script files, and you're correct in pointing out that if users want to run a PowerShell script file, they must first use the Set-ExecutionPolicy cmdlet to set the execution policy. They can find information about the cmdlet by retrieving its Help file in PowerShell.

In later installments, I do plan to discuss how to run script files. We've divided the information into two series: PowerShell 101 and PowerShell 201. In addition to the topics I cover in the first and second installments (February and March 2008), the PowerShell 101 series will cover how to use operators and work with expressions, how to work with string values in your commands, how to use variables in your commands, and how to work with PowerShell providers to access various data stores. The PowerShell 201 series will cover how to implement flow control in your commands, how to work with data types, how to create and use functions, and how to persist PowerShell scripts through script files and profiles. I hope this series provides a solid foundation for using PowerShell.

—Robert Sheldon

Hyper-V vs. ESX

I want to thank Michael Otey for his look at Microsoft's Hyper-V thin hypervisor ("A First Look at Windows Server 2008 Hyper-V," February 2008, InstantDoc ID 97857). In principle, I agree that Hyper-V might drive adoption of Windows Server 2008. However, I disagree that Hyper-V levels the playing field with VMware's ESX Server. Here's why:

- VirtualCenter—Michael mentions System Center for managing Hyper-V. Although I agree that managing some aspects of ESX can be tricky, VirtualCenter simplifies the management of resources, performance, migrations, and virtual machine (VM) creation across your enterprise.
- HA and Distributed Resource Schedule (DRS) clusters—These VirtualCenter features let you move VMs among ESX servers. DRS is particularly useful because you can create rules that prevent redundant VMs from being hosted on the same ESX server. (Think domain controllers—DCs—or MSCS nodes.) Keep in mind that all the migrations occur live.
- Community—The VMware Technology Network (VMTN) is a vibrant, technically strong community that solves many challenges of using ESX autonomously.
- VMFS and .vmdk files—VMware's VMFS file system and .vmdk are far more elegant than the partition model that Hyper-V uses. Instead of requiring individual partitions for guests, you

EDITOR'S NOTE

Windows IT Pro welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

create large VMFS volumes. Several VMs can happily coexist on a VMFS volume. The virtual disk file (.vmdk) is essentially a virtual partition in a file.

Hyper-V might be fine for the SMB, but in my opinion ESX is still the only real contender from an enterprise perspective. That said, Michael's article got me excited by Microsoft's progress because I'm sure it'll spark innovation from VMware. The recently released ESX 3.5 is early evidence of VMware's initiative. Several new experimental features have been integrated to improve the overall ESX experience.

—Brent McCraney

Thank you for writing. Watch for Michael Otey's virtualization shootout, in which he compares Hyper-V with ESX Server 3.5, in the June issue.

—Amy Eisenberg

Create SPF Records Automatically

I read the very useful Reader to Reader article by Nolan Garrett and Jeff Jones about SPF records ("Fighting Spam and Phishing with SPF," InstantDoc ID 98034, March 2008). I just want to add that Microsoft provides an easy Sender ID Framework SPF Record Wizard that automates the procedure of creating SPF records. You can find it at www.microsoft.com/mscorp/safety/content/technologies/senderid/wizard.

—Apostolos Fotakelis

InstantDoc ID 98587

Following the release of Windows Server 2008, Microsoft provided the following statement.

Sean Deuby's "Volume Activation in Server 2008" (InstantDoc ID 98153) offers a candid view of the activation technology manifested in Windows Vista and Windows Server 2008 and how IT pros should approach it while deploying Windows Vista and/or Windows Server 2008. However, there are few errors in the article. For example, the author mistakenly states that when the KMS client fails to renew with the KMS host within 80 days, an unusable system results until reactivation. In such scenarios, the resulting experience is notification, not any sort of limited use of the previously activated system. Additional changes to take note of: The initial grace period for Windows Server 2008 is 60 days, the default port for KMS location discovery is 1688, the KMS activation threshold is cumulative between Windows Vista and Windows Server 2008, and reactivating the system that has been previously activated using the MAK key is possible and results in "number of activations used" incremented by one. Updated prescriptive guidance is available at www.microsoft.com/technet/volumeactivation.

What You Need to Know About ...

Microsoft Response Point

Most small businesses don't have the resources or time to deal with traditional PBX phone systems, and even centrally managed solutions like Centrex can be expensive and complex. On the flip side, modern software-based solutions like Microsoft Unified Communications Server (UCS) are hugely expensive and complex, and UCS requires onsite technical staff.

You might expect Microsoft to simply retool its UCS product line and offer a "lite" version, perhaps integrated into a specific Small Business Server (SBS) SKU, which would provide small companies with IP-based phone capabilities. But such a thing would still be needlessly complex. What small businesses really want is a phone system that looks, acts, and works exactly like a phone. That's where Response Point comes in.

What's the Big Deal with Response Point?

Using a blend of old and new technologies—traditional looking, user-friendly phone handsets that utilize Windows XP Embedded (XPe) technology—Response Point is essentially a small business phone system in a box. OK, it's a big box: Response Point comes from one of several Microsoft hardware partners and is typically sold in a starter kit that combines a base station with five phones. You can add more phones at any time, and, now with the release of Response Point SP1, get seamless VoIP integration.

Response Point is designed for businesses with 50 or fewer employees, and unlike software-based solutions, if you ever do outgrow the system, you're going to have to move on. That said, the Response Point experience is as familiar as any phone system found at larger businesses, and unlike those systems, it's simple to configure. There's no management per se: You'll only need to occasionally change extensions, add and remove phones, and so on as the needs of your office change.

What's Really Happening Here?

A typical Response Point solution includes an XPe base station and several phones. The base station resembles a large network router, and it connects to both the business's external phone line and the internal wired network. There are no moving parts: The storage in the base station is all solid state and will work silently in a closet or other area for years. The only software installation required is for the management software: This will need to be installed on at least one PC in the office so that you can configure the phone systems and each phone. Installing individual phones requires only that you connect them to an Ethernet jack. Aside from the different wire, the phones are otherwise identical to standard office phones. Various types are available, from simple handsets to more full-featured designed with more capabilities.

An office manager or worker should be able to get the new phone system up and running in literally minutes: I was able to configure a review unit in about 15 minutes. If your office is truly incapable of figuring this out, various Microsoft partners sell and configure Response Point for a small charge. Unlike some subscription-based services, however, it's unlikely you'll need anyone to monitor the system regularly.

When you plug a phone into the office network, it shows up in the administrative console. From there, you can tie the phone to a specific employee ("Mary"), or, intuitively, to a location in the office ("receptionist," "meeting room"). This has several advantages: If you

Paul Thurrott (thurrott@windowsitpro.com) is the news editor for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.winiformant.com).

What You Need to Know About...

Windows Vista SP1 Scheduling

Microsoft took an unconventional path to deliver its first major update to Windows Vista. First, the company spent several months trying to convince its customers that SP1 didn't even exist and might not be released. Then, in late summer 2007, the company came clean and released a series of SP1 beta versions. Throughout the end of 2007, the SP1 feature set changed as well: Microsoft made changes to Vista's instant search functionality to appease Google, then changed the Windows Genuine Advantage (WGA) antipiracy functionality dramatically. Finally, in early February 2008, Microsoft completed SP1. But there was a final bit of controversy: Microsoft would release SP1 to customers under a strange, tiered schedule. Here's what you need to know about the release schedule for Vista SP1.

What's Going On with Vista SP1's Schedule?

The mystery behind SP1's release schedule is tied to the way Windows is developed internally: Both Vista SP1 and Windows Server 2008 were developed in lockstep and utilize the same code base. (So much so, that the first Server 2008 service pack will actually be SP2, allowing Microsoft to work on the next major updates to both Vista and Windows Server simultaneously.)

Server 2008 has been in development for almost five years, and Microsoft spent much of the past year refining that release and improving its overall quality, reliability, and fit and finish.

Because they were developed together, Microsoft wanted to freeze development of Vista SP1 and Server 2008 at the same time. Vista SP1, however, had a much shorter development cycle than Server 2008, and late in the beta process, Microsoft's beta testers discovered some device driver incompatibilities with SP1 that could cause customer PCs, which had worked fine with the original version of Vista, to not recognize particular hardware devices. While it's unlikely that Microsoft will ever identify these devices specifically, my sources tell me that they are mostly networking oriented, and some are quite common.

To meet its arbitrary internal schedule, Microsoft released Vista SP1 to manufacturing alongside Server 2008 in early February 2008 but announced that it wouldn't immediately release the update to the public, citing the mysterious device driver issues as the cause. Instead, the company said it would roll out SP1 to customers on a tiered schedule over the first half of 2008.

What's the Exact Schedule?

Depending on how you choose to acquire Vista SP1, it should be available by the time you read this, or shortly thereafter. Here's the schedule:

Beta testers. Those who participated in the Vista SP1 beta test received the code a week after Microsoft shipped SP1, on February 8, 2008.

Volume license customers. Corporate users who participate in Microsoft's volume licensing programs received DVDs with an integrated version of Vista and SP1 in February. MSDN and TechNet subscribers. Developers and IT pros who subscribe to Microsoft's MSDN and TechNet Plus services received Vista SP1 before the end of February.

Standalone downloads. The standalone versions of the SP1 code will be made available on Microsoft.com in mid-March.

Windows Update. Those wishing to upgrade an existing Vista-based PC to SP1 via Windows Update will be able to do so in mid-March 2008 or in mid-April 2008, depending on their hardware configuration. If you enable Windows Update to automatically download updates and do not have any of the affected hardware installed, you will automatically receive SP1 in mid-March. Otherwise, it will be mid-April.

Preinstalled on new PCs. New PCs with Vista SP1 will appear on store shelves "in the coming months," according to Microsoft. PC makers began receiving the SP1 code in early February, but it will likely take at least a few months before new PCs based on that code hit the market. My guess is that Vista SP1 will be the default installation on new PCs by mid-year.

Retail copies of Vista SP1. If you'd like to purchase a boxed copy of Vista SP1, those versions of the system will replace the initial Vista release on store shelves "in the coming months" as well. Again, this process should be completed by mid-year.

International users. The version of SP1 that was completed on February 4, 2008 included only the English, French, Spanish, German, and Japanese language versions of the update. The remaining languages that Microsoft supports will be released to manufacturing in April and ship worldwide after that, the software giant says.

Recommendations

There's nothing like a final bit of silliness to cast yet another cloud of suspicion over a release that should have been a slam dunk for businesses that were waiting on SP1 to deploy Vista. That said, Microsoft says it won't need to make any changes to the SP1 code to correct the device driver incompatibilities. Pedantic arguments over what constitutes "final" notwithstanding, SP1 is a high-quality release that does remove some important deployment blockers and improves Vista's overall performance, reliability, and security. My advice here is simple: You should be installing SP1 as soon as possible. And if you haven't yet deployed Vista, your excuses are running out. SP1 puts Vista over the top, despite the tiered release schedule.

InstantDoc ID 98443

rotate through part-time help regularly, it might not make sense for particular phones to be tied to particular employees. You can also create groups ("Sales") so that specific calls can be routed appropriately to more than one phone. No matter how you provision the phones, it's simple to change their configuration at any time through the console.

One of Response Point's best features, and the one that's transformed my home office into a multinational corporation as far as anyone who calls is concerned, is a speech-based automated receptionist. Callers can navigate through the phone system using a pleasant and surprisingly useful voice-recognition feature that routes calls accordingly. Naturally, it also supports voice mail (with up to 1000 minutes of messages per base station) and call routing and forwarding. So if your best salesperson is out on a call, important calls can be forwarded to him or her—or not: The system is very flexible.

Pricing is also in line with the budget of a typical SMB. Microsoft partners including Quanta, D-Link, and Aastra Technologies are now selling starter packages (one base station and four to five phones) for about $2500. Additional wired phones cost about $160 each, and Microsoft tells me that wireless versions will arrive sometime this year. Response Point comes with no hidden fees. Once you purchase the hardware, it's yours to use for as long as you'd like with no additional costs. Take that, UCS.

Recommendations

Response Point is a wonderful and capable system for small businesses that don't have a dedicated tech staff and aren't likely to keep a Microsoft partner on retainer to manage a more complex system. However, you won't be able to grow with Response Point past the 50-user mark, at which point it might be time to consider more complex systems and the resulting support requirements. But I'm surprised and heartened to see that Microsoft can make such a solution: Response Point doesn't require Windows Server, Active Directory, Exchange, or any of the software giant's other hugely successful but complex enterprise products. It is, in other words, perfect for the typical small business.

InstantDoc ID 98444

Find Out Who Is Leaving Files Open at Night

Our office uses EVault for nightly backups of our file servers. When reviewing EVault's backup logs, I noticed that some files weren't being backed up. Although the backup logs noted that these files weren't being backed up because the files were open, the logs didn't state which users or processes opened the files.

I knew that I could log on to the file server when a backup was being performed to try to determine which users or processes are leaving files open. However, the backups are performed at night, so I did some digging to try to find another option.

I found that Windows Server 2003 and later server OSs have a built-in command-line tool named Openfiles that you can use to display open files. (This tool is also included in Windows XP and later client OSs.) By using a batch file similar to the one I used to schedule Windows 2003's defragmenter (see "Automate the Windows 2003 Defragmenter Without Paying Extra," May 2007, InstantDoc ID 95487), I found that I could send the output from Openfiles to a text file to obtain a snapshot of which users and processes had files opened when a backup was performed.

The batch file, chk-open-f.bat, runs Openfiles and sends its output to a text file named Openfiles-1.txt. The Openfiles-1.txt file contains four columns—"ID," "Accessed By," "Type," and "Open File (Path\executable)"—filled with information that you can use to determine which users or processes have files opened. I run chk-open-f.bat every evening before EVault runs.

As Listing 1 shows, chk-open-f.bat begins by performing a backup of its own. Because I wanted to run this batch file every day, I needed a way to preserve Openfiles' output over the course of a weekend. In other words, I needed to see the Openfiles output from Friday night, Saturday night, and Sunday night come Monday morning. So, chk-open-f.bat first backs up previous text files for up to three days, so you can have a total of four text files. (Only after the first four nights will you see all four text files.) If you want to back up more or fewer text files, you can easily edit the code. Each text file is created on the root of the file server's E drive. You can change this location if needed by editing the path.

After backing up the text files, chk-open-f.bat deletes any existing Openfiles-1.txt file so that Openfiles-1.txt contains only the Openfiles output from the current night's run. (Otherwise, the Openfiles output from the current night's run would be appended to the existing data in that file.) Finally, chk-open-f.bat runs Openfiles, sending its output and a date and time stamp to Openfiles-1.txt.

To use chk-open-f.bat, customize it if needed, make it a read-only hidden file for security purposes, and place it on the root drive of the server you want to monitor for open files. Then, schedule chk-open-f.bat to run. You can use Task Scheduler or the AT command to schedule chk-open-f.bat. However, I don't use Task Scheduler to run chk-open-f.bat because I'm running Openfiles on a file server, and usually no one is logged on to that server. For a scheduled task to run when no one is logged on, you need to supply a username and password that has admin rights. I'd rather not specify this information just in case someone learns how to hack into Task Scheduler to obtain admin passwords. Instead, I use the AT command in a scheduler script to run chk-open-f.bat like a service—that is, without a logon account. Plus, if you use this approach, the Openfiles process runs in the background, which prevents windows from popping up.

Listing 2 shows a sample scheduler script named Set.bat. As currently set up, Set.bat schedules chk-open-f.bat to run every night at 6 P.M. If you use this script, you'll need to customize the days and time when you want chk-open-f.bat to run. When doing so, be careful not to change the spacing in the AT command; changing the spacing might cause chk-open-f.bat to be scheduled incorrectly. Also, make sure the path to chk-open-f.bat points to proper location. After you customize Set.bat, save it to the same location as chk-open-f.bat on the server, then execute it. You won't need set.bat anymore, so you can delete it from the server.

I use chk-open-f.bat and set.bat on Windows 2003 servers. I haven't tried them on any other OSs. You can download these scripts by going to www.windowsitpro.com, entering 98553 in the InstantDoc ID box, clicking Go, then clicking the *Download the Code Here* button.

Daniel L. Gillard

EDITOR'S NOTE

Share your Windows discoveries, comments, solutions to problems, and experiences with products and reach out to other Windows IT Pro readers (including Microsoft). Email your contributions to r2r@windowsitpro.com. Please include your phone number. We edit submissions for style, grammar, and length. If we print your submission, you'll get $100. Submissions and listings are available online at www.windowsitpro.com. Enter the InstantDoc ID number in the InstantDoc ID text box.

Listing 1: Chk-open-f.bat

```
@Echo Off
Copy e:\openfiles-3.txt e:\Openfiles-4.txt /y
Copy e:\openfiles-2.txt e:\Openfiles-3.txt /y
Copy e:\openfiles-1.txt e:\Openfiles-2.txt /y
Del e:\openfiles-1.txt
Echo Checked %date%,%time% >>e:\OpenFiles-1.txt
Echo ------------------------------------ >>e:\
  OpenFiles-1.txt
openfiles >> e:\Openfiles-1.txt
Echo ------------------------------------ >>e:\
  OpenFiles-1.txt
```

Listing 2: Set.bat

```
At 06:00pm /every:M,T,W,TH,F,S,SU e:\chk-open-f.bat
```

Chk-open-f.bat is a simple but useful batch file if you're having problems with files being left open when they shouldn't be. Using chk-open-f.bat, I can identify which users are leaving files open at night so that I can talk with them individually about the problem. This approach is much more effective than sending all users a generic message about not leaving files open and hoping the offenders realize that the message is targeted at them.

—Daniel L. Gillard, LAN administrator, Rogers Media

InstantDoc ID 98553

Enable Remote Desktop Over a Network On Demand

Remote Desktop is one of the most important support tools for administrators of remote computers, but it's often not enabled on desktop systems when you need it. If you know what you're looking for, however, it's fairly easy to enable it remotely on a network.

The Remote Desktop service on desktop systems is always running, even if remote access is disabled. Whether remote connections are enabled or disabled is controlled by a value named fDenyTSConnections under the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Terminal Server. By default, this value is 1 (which is disabled). When you use the Remote tab in the System Properties dialog box to enable Remote Desktop access (shown in Figure 1 for a computer running Windows Vista), this value changes to 0.

Because you can enable remote connections through the registry, you have a wide range of options for modifying Remote Desktop access. Essentially, any technique for accessing a registry value will work, including the following:

- You can open regedit, connect to the remote computer, then change the value.
- You can have a PC import a .reg file on startup.
- You can use Windows Management Instrumentation's (WMI's) System Registry provider to make the change from a script.
- You can use reg.exe to make the change from the command line or a script.

Many administrators seem to need to enable Remote Desktop access on the fly. It's also usually a good practice to turn the remote connection back off when you're done using it. So, to simplify the process of checking whether Remote Desktop is enabled, enabling it if it's not, and disabling it after you're done using it, I wrote three one-line .cmd scripts: CheckRD.cmd, RDOn.cmd, and RDOff.cmd.

To check whether a computer has Remote Desktop enabled, you can run the CheckRD.cmd script in Listing 3. To download it and the other scripts, go to www.windowsitpro.com, enter 98551 in the InstantDoc ID box, click Go, then click the *Download the Code Here* button. CheckRD.cmd and the other two scripts take a single command-line argument: the name of the remote computer to connect to. For example, if you want to use CheckRD.cmd to see whether Remote Desktop is enabled on the remote computer named Client07, you'd run the command

```
CheckRD Client07
```

You can also use this script to check the setting on the local system. Just specify a period (.) as the name of the host computer. No matter whether you check a local or remote system, the script returns the value of fDenyTSConnections. If Remote Desktop is disabled, 0x1 is returned. If it's enabled, 0x0 is returned.

To enable Remote Desktop, you can run RDOn.cmd, which Listing 4 shows. You need to run this script from an account that has sufficient privileges to change the remote machine's registry. To enable access on the Client07 computer, you'd run the script with the command

```
RDOn Client07
```

You'll be able to immediately make a remote connection.

If you need to disable Remote Desktop after you're done using it, you can run RDOff.cmd in Listing 5. For example, to turn Remote Desktop off on the Client07 computer, you'd use the command

```
RDOff Client07
```

As you can see, CheckRD.cmd, RDOn.cmd, and RDOff.cmd are simple scripts. But don't let their simplicity deceive you. These scripts let you easily enable and disable a computer's remote connection on demand.

—Alex K. Angelopoulos, Senior Network Engineer

InstantDoc ID 98551

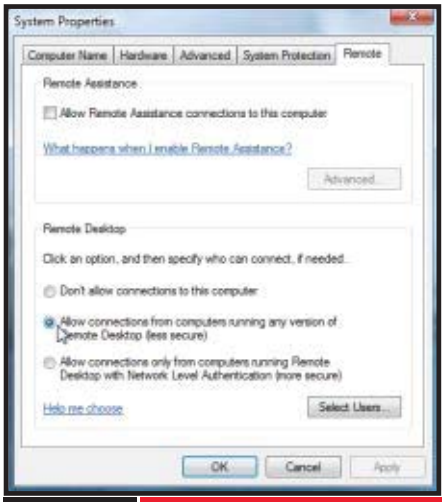

Figure 1: Using the Remote tab in the System Properties dialog box to enable Remote Desktop access

Listing 3: CheckRD.cmd

```
@reg query "\\%1\HKLM\SYSTEM\CurrentControlSet\Control\Terminal Server"
  /v fDenyTSConnections
```

Listing 4: RDOn.cmd

```
@reg add "\\%1\HKLM\SYSTEM\CurrentControlSet\Control\Terminal Server"
  /v fDenyTSConnections /t REG_DWORD /d 0 /f
```

Listing 5: RDOff.cmd

```
@reg add \\%1\HKLM\SYSTEM\CurrentControlSet\Control\Terminal Server
  /v fDenyTSConnections /t REG_DWORD /d 1 /
```

… Tired of Nursing Your Exchange Server?

#1 BEST SELLER!

Anyone who has given birth to an Exchange network knows it can get sick and needs some nursing to stay healthy. In fact, 72% of Exchange Administrators surveyed* have "experienced" an Exchange disaster (feels like the flu)—usually from improper feeding and care.

Like many databases, constant adding and deleting can corrupt an Exchange data file so it eventually turns sour. Replicating, archiving and backing up the data doesn't stop the stink—it just stores it. You've got to…

Fix the Problem

You may have tried the free utilities to fix Exchange. While they help, they are too tedious, time consuming and lightweight to keep your Exchange baby healthy. You've tried the milk, now try some meat!

Pamper Yourself with GOexchange

It's time to try GOexchange, from Lucid8, the #1 best-selling automated disaster prevention and optimization software for Microsoft Exchange 5.5, 2000, 2003 and 2007. As the mother of all Exchange tools, GOexchange helps prevent disasters, repair problems, improves performance, and saves you a lot of time.

"Without routine maintenance, decreasing performance, increased warnings and errors accumulate and database fragmentation transpires, leading to Exchange disasters."

Gartner

Prevent Hiccups

GOexchange removes errors, warnings and inconsistencies within the database—before major corruption makes the database fail.

"GOexchange corrected 2,264 errors and 26 warnings."

Paul Ramos, Director IT

Run, Don't Crawl

In addition to fixing the database, GOexchange removes sluggishness and improves performance by re-indexing and defragmenting the database to permanently remove white space and deleted items. The end result is increased performance and stability with a compact efficient database that's 31 to 55% smaller! Combine this with archiving and the database is up to 91% smaller—making it much quicker to backup.

"..our information stores were reduced by 45-50%."

Dale Huitt, Systems Lead

Automated Babysitter

First, GOexchange is easy to setup and use. Twenty minutes—that's all it takes to get your server up and running. Just schedule it, and walk away!

The software notifies the users, validates the database, runs the backup, conducts a comprehensive system analysis and diagnostics, logs the errors, and notifies you if it discovers a "stop" error—then it repairs and defragments the database, generates a thorough report and schedules the next event.

You can do some of this work yourself, but why waste time doing repetitive maintenance, when GOexchange can do it for you—faster and more effectively than doing it by hand.

Created By

Lucid

Solutions Inspiring Confidence

"Life before GOexchange…was an absolute nightmare, late nights, long weekends and upset users."

Marty Grogan, CTO

Stop The Crying

Why not call now, or visit our resource site and learn how to reduce the risk, and avoid the pain. Protect your exchange data, maximize performance, and spend a weekend at home—*instead of babysitting Exchange.*

Special Offer

- Free Software for analysis of your Exchange server!
- Free White Paper—"Basic Feeding of Your Exchange Server."
- Free Essential Guide to Exchange Preventative Maintenance

Go to: www.Lucid8.com/GoITPro
Call 425.456.8474
E-mail: Sales@Lucid8.com

Special Operations Software™

Powering Your Active Directory through Innovation and Simplicity!

Specops Command

PowerShell remoting through Group Policy

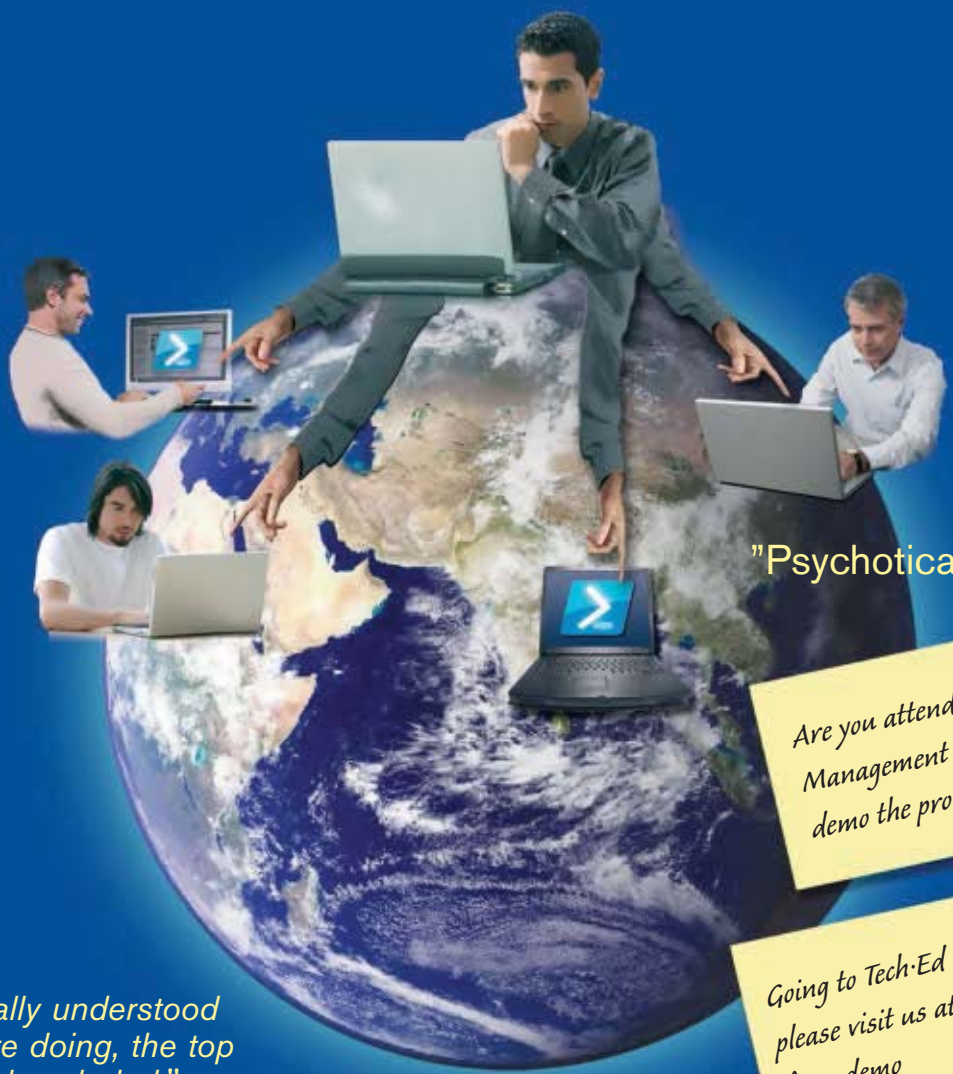

"Psychotically Powerful"

Are you attending the Microsoft Management Summit, let us demo the product at booth 517

Going to Tech·Ed IT Professionals, please visit us at booth 415 for a demo

"...when I finally understood what they were doing, the top of my head exploded."

Jeffrey Snover
Windows Management Partner Architect

Read more about Jeffrey's impressions of Specops Command at the MSDN PowerShell blog:
http://blogs.msdn.com/PowerShell/

Specops Command ™
We bring you the future of scripting, today!

Microsoft GOLD CERTIFIED Partner
Security Solutions
ISV/Software Solutions

– For more information about Specops Command and how to download your FREE limited version please go to:
http://www.specopssoft.com/powershell

EDITOR'S NOTE: Send new product announcements to products@windowsitpro.com.

Network Management

Monitor Network Availability and Performance

Network Instruments has announced the availability of **Link Analyst 5.1**, the latest version of their network infrastructure management and performance monitoring application. According to Network Instruments, this latest release uses SNMP and WMI data to identify network performance issues by monitoring the availability of attached network devices. It also can be used to keep tabs on network performance by departments, groups and by business processes. LinkAnalyst 5 enables the monitoring of WAN, LAN, wireless, Fibre Channel, gigabit and 10 GbE networks. Link Analyst 5 is available now, and pricing begins at $2,495. For more information, contact Network Instruments at 952-358-3800 or visit www.networkinstruments.com.

Infrastructure Security

Certificate Management

ChosenSecurity Inc.has announced **TrustCenter Enterprise ID QuickStart** (TC EID QuickStart), a new on-demand certificate management service for small and mid-sized businesses. TC EID QuickStart gives SMBs the ability to add security features such as digital signatures and encrypted email to their IT environments. The service also provides management features that allow IT pros to create, edit, suspend, and delete user and device security profiles. Pricing for TC EID QuickStart begins at $20 per user, per year. For more information, contact Chosen Security at 866-468-2180 or visit www.chosensecurity.com.

Product Spotlight

Hardware

Mobile Thin Clients

Wyse Technology used VMworld Europe 2008 to unveil two new mobile thin clients: the **Wyse X90L** and **X90Le**. Both devices ship with Windows XP Embedded, Gigabit Ethernet, 802.11b/g/n/ wireless, Bluetooth 2.0 support, 15.4" TFT screens, and smart card reader with Citrix Password Manager security software. Solid state memory is used for storage, and the lack of moving parts helps reduce power consumption. According to Wyse, these new thin clients will ship with several virtualization-friendly features. "Customers have told us that it makes no sense to access a virtualized PC in the datacenter with another PC on the desktop. They want a secure device that delivers a rich experience, and they want to enable users to take that experience with them from their office, to the conference room, even moving off campus to their home or other locations," said Jeff McNaught, chief marketing officer for Wyse Technology, in a statement announcing the new products. "We went well beyond the idea of stripping down a notebook PC, to deliver a thoughtful, future-proofed combination of performance, peripheral support, and connectivity, offering Wi-Fi b, g and even n and Bluetooth 2.0." For more information, contact Wyse at 408-473-1200 or visit www.wyse.com.

Email Security

SPAM Protection

Preventing email spam is the focus of Abaca Technology's **Email Protection Gateway.** According to Abaca, the appliance **relies** on patent-pending technology—dubbed ReceiverNet—that the vendor claims has a 99 percent spam prevention rate. ReceiverNet evaluates email based on the idea that legitimate users receive less spam, and that email accounts that receive lots of spam tend to send more spam. RecieverNet then rates incoming messages accordingly, and also blocks most malware, viruses, and phishing attacks. For more information, contact Abaca Technology Corporation at 408-571-6400 or visit www.abaca.com.

Virtualization

VM Security

Security for virtualized environments is becoming increasingly important, so VMware has introduced **VMware VMsafe**, a new virtualization security product. According to VMware, this

new product "protects applications running in virtual machines in ways previously not possible in physical environments." In a news release announcing the product, VMware explained that VMsafe works with the VMware hypervisor to increase security against threats from malware, keyloggers, viruses, and other malicious software. A VMsafe API will allow vendors to develop security products that integrate with VMsafe; VMware announced that is working with Check Point, McAfee, and Symantec on new security solutions that support VMsafe technology. For more information, contact VMware at 877-486-9273 or visit www.vmware.com.

Network Management

Manage Network Inventory

Developer Geert Moernaut has announced the release of **Lansweeper 3.0**, the latest version of his popular freeware network asset detection utility. IT pros can use Lansweeper 3.0 to make comprehensive inventories of the hardware and software located on their network infrastructure. The application is server-based, and doesn't require any clients to be installed on any workstations in your network. Network scanning is performed using registry, file shares, and WMI information. More than 75 reporting templates are included, and data can be exported to Microsoft Excel for further analysis. For more information, contact Geert at geert@moernaut.com or visit www.lansweeper.com.

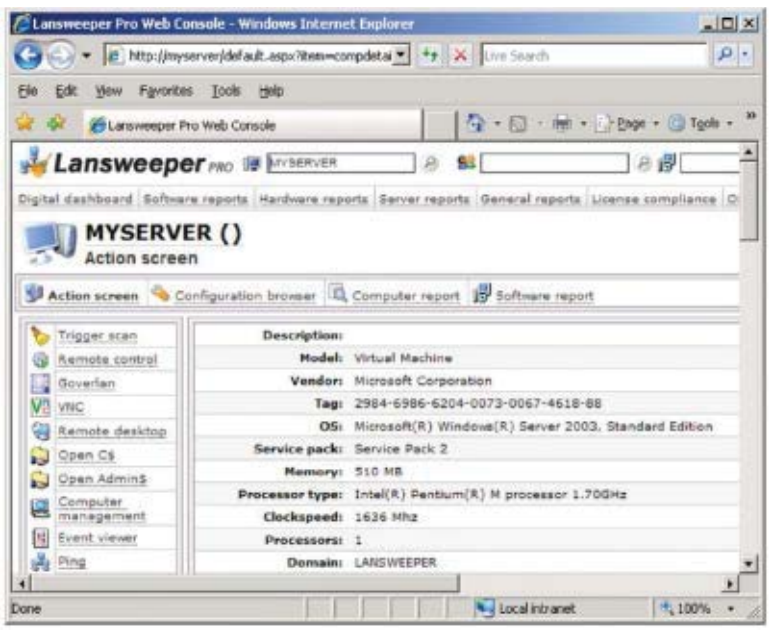

Hardware

Data Center Infrastructure

Brocade has unveiled the **Brocade DCX Backbone**, the first in a series of new data center networking products. This new product includes support for up to 896 ports of 8 Gbit/sec Fibre Channel, and also includes reduced power consumption features. According to Brocade, this new range of products relies on adaptive network services to dynamically distribute resources between servers (both physical and virtual) and networked storage. Pricing for the Brocade DCX Backbone begins at $180,000. For more information, contact Brocade at 408-333-8000 or visit www.brocade.com.

InstantDoc ID 98459

_INFRASTRUCTURE LOG

_DAY 85: Woke up in a desert. Our data center is so hot it's playing tricks on our minds. We have to do something about these energy costs. But how?

_Maybe that sphinx over there has an answer.

_DAY 86: We need IBM! Their services can help us plan and build a more energy-efficient data center. A virtualized IT environment can improve our server and storage utilization, while their power management capabilities help us actively manage our power usage.[1] And IBM's advanced cooling solutions can make our data center cucumber cool.

Tivoli®

Watch a Webcast on data center energy efficiency at:
IBM.COM/**TAKEBACKCONTROL**/ENERGY

[1]Available on select models.

Insights from the industry

NetPro's Exchange Solution Route

"The massive number of e-mails that go through an organization each day makes managing the Exchange system increasingly complex and time-consuming. That goes without saying. But keeping up with system upgrades and configuration changes is a whole other story," says NetPro's Brad Hibbert, vice president of strategy. NetPro believes its new offerings can help ease Exchange admins' pain.

Known for its Active Directory (AD) expertise, NetPro is making a major move into the Microsoft Exchange market. NetPro's Exchange solutions, some of which are already shipping and others which will ship later this year, variously help you manage, migrate to, analyze, monitor, recover, and archive in your Exchange universe.

- NetControl for Exchange, shipped last February, leverages NetPro's service-oriented architecture to help admins nail down the who, what, where, when, and why of Exchange users.
- NetMigrate for Exchange, which also shipped last February, offers help in migrating to Exchange 2007, with minimal impact on users.
- ChangeAuditor for Exchange 4.5, shipped in March, builds on NetPro's existing ChangeAuditor for Exchange product, offering tracking and monitoring functionality for Exchange 2007.
- NetControl for Exchange Archive and NetControl for Exchange Message-Level Recovery are due out later this year and add pieces to the Exchange management puzzle with archiving and recovery functions.

Hibbert cited the increasing popularity of Exchange Server 2007 with customers. "A lot of folks are moving to Exchange 2007 because of compliance reasons," Hibbert says. "Our existing customers asked for it, and it means we can leverage our technology and distribution." (Pricing starts at $6 per seat for NetControl, $18 per seat for NetMigrate, plus a maintenance percentage. The solutions will also be available as a suite.)

—Caroline Marwitz

InstantDoc ID 98242

The Power of 3: ScriptLogic Incident Management Solution

ScriptLogic Incident Management Solution came out in February and addresses three areas IT pros need help with: responding to problems, managing problems, and preventing problems. It bundles tried-and-true tools Desktop Authority Remote Management Gateway and ScriptLogic BridgeTrak with newcomer Desktop Authority Password Self Service, in a solution that seems indicative of a general trend that empowers users to become part of the answer instead of being part of the problem.

ScriptLogic, a subsidiary of Quest Software, acquired BridgeTrak's Help desk solution, the second part of the ScriptLogic Incident Management trio. BridgeTrak helps admins track and manage problems and includes a knowledge base so users can find solutions themselves. Cavalancia says that Forrester estimates the knowledge base concept alone results in a 15 percent reduction in user calls to the Help desk.

Lastly, Desktop Authority Remote Management Gateway lets admins remotely control and manage user desktops whether users are on the corporate network, the VPN, or totally off. Users can continue working while their system is under remote management. Cavalancia says the integration of the three products is a natural fit. The cost is based on the number of admins (for the Help desk piece, a minimum of five), and the number of users (for the remote management piece). One hundred users can run a little over $5,600.

—Caroline Marwitz

InstantDoc ID 98235

Dell Makes SaaS Move, Deals for MessageOne

An interesting new union is taking place. Perhaps you've heard about it: Dell has struck a deal to acquire MessageOne, well-known provider of email management services such as archiving and security. The question I think many people are asking is "Who is this good for?" Well, it could be you.

This deal potentially puts Dell in a position to compete in the growing Software as a Service (SaaS) market with the likes of Google (Google Apps) and Microsoft (Exchange Hosted Services). And remember, competition is good for you, the consumer. Dell has been moving in the SaaS direction and recently announced its ProSupport service, to which the MessageOne services are likely to be added. Dell should have some great new features to help promote its hardware sales.

MessageOne's email continuity services are well-established, providing continuous uptime and ease of implementation. As Paul D'Arcy, vice president of marketing for MessageOne, told me, "An important goal of the acquisition is to bring these products to a much broader audience. As MessageOne becomes part of Dell, we will have incredible resources to accelerate innovation and to broadly expand global delivery of these products." That sounds good, doesn't it? So maybe everybody wins with this deal.

Coincidentally (or not?), MessageOne recently announced the availability of its EMS Rapid Archive, which helps organizations quickly set up an email archive (in as little as a day) and manage retention policies, e-discovery, and litigation holds. This solution lets you start your archive with any number of mailboxes and grow as your needs change. And with zero onsite maintenance by your IT staff, the cost benefit of such a solution is clearly evident. EMS Rapid Archive integrates with MessageOne's EMS Email Continuity and EMS Email Security.

—Brian Winstead

InstantDoc ID 98276

ULTRABAC™
SOFTWARE

Reliability — Catch phrase or reality?

Hot backups. Business continuity. Continuous data protection.

These and other buzzwords have been generated by the technology industry to get your attention. But what do these terms mean to you? To UltraBac Software, they are another way of explaining the goal at the center of our business: *perfect uptime*. We invest the time and resources to offer solutions that ensure you have the fastest, most reliable access to your stored data, avoiding costly downtime waiting for failed machines to recover.

Introducing Continuous Image Protection (CIP). UltraBac Software is so excited about our new technology that aids in this goal, we're announcing CIP before its planned release. CIP is a form of *continuous data protection (CDP)* with a new innovation: it automatically backs up each sector on a disk as it is changed, unlike standard image backups which run only on a periodic basis. With CIP, your image backup never stops – so a system can be brought back to a point-in-time, rather than restoring a static image that could be up to 23 hours old.

UltraBac Software's sole mission is data protection. So when we advertize *product reliability*, *innovative features and functions*, and *top-notch support* they are not simply catch phrases we use, but rather our commitment to you and your business.

UltraBac – Innovative software from a reliable company.

25 YEARS

BACKUP AND DISASTER RECOVERY SOFTWARE FOR PEOPLE WHO MEAN BUSINESS > WWW.ULTRABAC.COM

IronKey IGB Enterprise Special Edition

Editor's Note: Following is a summarized version of Jeff James' review of IronKey IGB Enterprise Special Edition. To read the full-length version of the article, go to www.windowsitpro.com and enter InstantDoc ID 98345.

Security threats to your IT infrastructure can come from just about anywhere, but one of the most troublesome areas of exposure is mobile devices—particularly, removable USB flash drives. Efficiently securing the data that they hold can be difficult. IronKey has come up with a solution to that problem with its **IronKey IGB Enterprise Special Edition**, a secure flash drive that comes loaded with security features. The device is available in three storage sizes and two product lines: A standard version includes secure Web-browsing features for individuals, and an enterprise variant omits some of those more consumer-focused applications. Both variants feature secure online backup functionality that can help keep your data secure (and recoverable).

To benefit from the IronKey's security features, you first need to initialize and configure the device's security settings. When you use the IronKey USB Drive for the first time, it prompts you for a password. After that password is established, the drive initializes, and the device asks you to create a secure IronKey online account, which you can use later to disable the device in the event of theft or loss, recover passwords, and back up the contents of your IronKey device online.

The IronKey USB Drive comes with a custom cryptographic processor called the IronKey Cryptochip that has its own password-guessing counter. If you (or some nefarious criminal type) make 10 incorrect password attempts while trying to access the device, the IronKey Cryptochip blocks further attempts and triggers a self-destruct sequence that permanently destroys the data on the USB drive.

As for physical protection, the IronKey USB Drive is enclosed in a sealed, waterproof metal case, and the open space inside the drive is filled with an epoxy-based compound. If someone tries to open the case, he or she will likely destroy the components in the process. I submerged the device under water, dropped it repeatedly, and even stepped on it a few times, and it always worked reliably when plugged into a USB port.

Is the IronKey USB Drive worth the extra money you'll spend, compared with the cost of standard USB drives? I can tell you that if you're one of those paranoid IT administrators (and you know who you are), this rugged, impressive piece of hardware could be just what you need, particularly if you work for the military or regularly handle valuable information.

InstantDoc ID 98345

—Jeff James

SUMMARY

IronKey IGB Enterprise Special Edition

PROS: Rugged physical design; useful online backup and safety features

CONS: Much more expensive than standard USB flash drives; self-destruct feature after 10 password failures might make the device an expensive prospect

RATING: ◆◆◆◆◆

PRICE: $79 for 1GB; $109 for 2GB; $149 for 4GB

RECOMMENDATION: If you're looking for the ultimate in USB flash drive security, the IronKey USB Drive is as good as it gets.

CONTACT: IronKey • 650-492-4055• www.ironkey.com

Paul's Picks

Summaries of in-depth product reviews on Paul Thurrott's SuperSite for Windows
www.winsupersite.com

Microsoft Office 2008 for Mac

PROS: Compatibility with latest Office document formats; true Universal application offering native compatibility on Intel Macs

CONS: Lackluster performance; no innovative features; no ribbon UI from Microsoft Office 2007 for Windows

RATING: ◆◆◆◇◇

RECOMMENDATION: Macintosh users have been waiting two years for Microsoft to ship a Universal version of its Office suite that offers native compatibility with Intel-based Macs. (Previous versions were written for the PowerPC and ran under an emulation mode.) Well, it's here and, frankly, it's not much of an improvement. Office 2008 for Mac isn't demonstrably faster than its predecessor, and its non-standard UI looks nothing like other Mac apps and doesn't offer any of the niceties of the Office 2007 ribbon from Windows. But you might find it worth using if you need compatibility with Microsoft's latest XML-based Office document formats or with Microsoft Exchange.

CONTACT: Microsoft • 800-426-9400 • www.microsoft.com

DISCUSSION: www.winsupersite.com/showcase/win2008_rcl.asp

Windows Live SkyDrive

PROS: Anywhere-anytime access to data; easy to share data with others

CONS: Limited storage

RATING: ◆◆◆◇◇

RECOMMENDATION: Windows Live SkyDrive, Microsoft's Web-based storage service, is a rare example of the software giant beating Google to market. The service is only so-so: It offers 5GB of free storage but no way to increase that storage allotment. And it feels pokey, especially with file uploads. However, SkyDrive makes it very easy to share files securely with others and to specify whether they can edit or just view and download files.

CONTACT: Microsoft • 800-426-9400 • www.microsoft.com

DISCUSSION: www.winsupersite.com/showcase/win2008_rcl.asp

InstantDoc ID 98452

Speed up Your Systems in Real Time

The 8 Essential Benefits of Automatic Defragmentation

Fragmentation is unavoidable. It wreaks havoc on hard disks, causing crashes, hangs and complete system failures.

Diskeeper 2008 eliminates fragmentation—automatically. It's the real-time solution to your performance and reliability problems. Diskeeper is absolutely indispensable. It speeds up boot times, makes applications launch faster and improves the efficiency of backups and anti-virus scans. Diskeeper's benefits have proven time and time again to be a vital part of system administration.

We asked 254 of our customers what were the essential benefits of using Diskeeper. This is what they had to say:

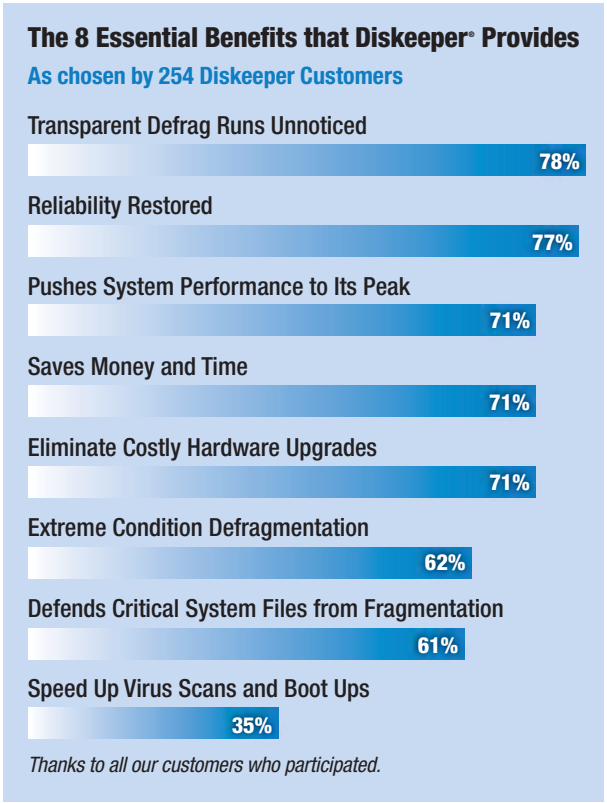

1. Pushes System Performance to Its Peak

"We had one machine that had a failing drive in a RAID 5 array and when we replaced that drive, performance improved by 300%. And then when I ran Diskeeper for a week, again it improved over 300%. A disk intensive process that was taking 1.5 hours is now taking 15 minutes."

2. Reliability Restored

"We use Microsoft® SQL Server®. We were receiving hundreds of messages per day in the log like this one: SQL Server has encountered 21 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [E:\mssql\data\…]

"We researched this error and found that it is usually caused by badly fragmented hard drives. While our drives are part of a large SAN solution, we were not totally convinced that this should be causing the problem. We downloaded a trial version of Diskeeper and after running it, all of these errors disappeared! We have purchased 5 copies of Diskeeper and we are installing them on all of our production databases with the expectation to never see this error again!"

3. Transparent Defrag Runs Unnoticed

"The server automatically defragments only when there are idle resources. No more worrying about when I can schedule defragmentation, no more worrying about if the defragmentation will cause performance issues. InvisiTasking™ has worked great for us on everything from file and print servers to SQL servers."

4. Defends Critical System Files from Fragmentation

"I have been using Diskeeper at my office on the 63 workstations and 4 servers over the last year. The addition of Frag Shield™ 2.0 eliminates the task of manually changing the MFT. In the past most of my MFTs needed adjustment. Now that this function is automatic, I don't have to manually check it."

5. Saves Money and Time

"Prior to installing Diskeeper, we were manually defragmenting. Some of the drives would take hours to defrag and within a few days we would need to defrag again. Installing Diskeeper basically paid for itself within a month by reducing off-hour salaries. Also the defragmented drives perform better and last longer. It's a no-brainer for production machines."

6. Speed Up Virus Scans and Boot Ups

"Diskeeper saves time in doing virus scans, backing up, indexing and searching the files. There are also faster download times for users because of the lower load on the defragmented RAID."

7. Extreme Condition Defragmentation

"One day our SQL Server came to a halt. I did everything: ran spyware software, deleted numerous .TMP files, ran Windows® update, etc. But nothing got the server to run. Then I installed and ran Diskeeper; I found that the hard drive was horribly fragmented. But after Diskeeper finished defragging the system, the server came up."

8. Eliminate Costly Hardware Upgrades

"We were looking at having to replace or upgrade some of the servers because they were so slow. Since the Diskeeper install, they are performing well enough that we are no longer looking at the upgrades and replacements."

Diskeeper is essential for maximum speed and reliability on networked systems. Accelerate your systems' performance. Restore reliability. **Try Diskeeper 2008 for free now!**

SPECIAL OFFER

Try it FREE for 45 days!

with InvisiTasking™ Diskeeper® 2008
Maximizing Performance and Reliability—*Automatically*™

Go to www.diskeeper.com/witp
(Note: Special 45-day trialware is only available at the above link)

Volume licensing and Government/Education discounts are available by calling 800-829-6468, code 4047.

3 ENTERPRISE DISK DEFRAGMENTERS

A low price and slow-but-successful performance determine our Editor's Choice

by Eric B. Rux

Monitors are getting thinner, CPUs are becoming faster, and software is getting easier to use. But one aspect of computing that remains the same is hard disk fragmentation. Our OSs—for various reasons—are failing to store files in contiguous disk space and are instead tossing parts here and there, filling in disk gaps willy-nilly. Our poor hard disks have to read portions of files scattered all over the platter, rather than reading files in smooth, continuous motions.

So, you need a defragmentation tool. Where do you start? We've selected three products for a comparative review—Diskeeper's **Diskeeper 2008 Server**, O&O Software's **O&O Defrag 10 Server Edition**, and Raxco's **PerfectDisk 2008 Server**—that should help you decide which best suits your environment.

All three companies also offer workstation editions (and even products for Exchange Server, SQL Server, and Windows Home Server), but I found few differences between these server and workstation versions. The primary goal of both is to simply defragment a computer's hard disk. However, the server versions in this review either come with built-in enterprise functionality or offer a separate add-on to help you centrally manage your defragmentation tasks.

Testing

To permit straightforward comparison of the products' features, I used VMware Server 1.0.4 to install each product on a virtual machine (VM). However, I also felt it was important to install each product on actual hardware to compare performance results, so I did that, too. To ensure that I compared the products fairly, I used disk-imaging software to capture a heavily fragmented hard disk with only 20 percent free space. In addition, I included one extremely fragmented file that was larger than the free space, as well as one heavily fragmented disk with a lot of free space. In just a few minutes, I could easily reproduce the fragmented drives for testing. Then, I used each product to run one manual defrag pass on the hard disk. You can find the results in Table 1, page 26.

Product Overview

All three products offer offline and online defragmentation. Offline defragmentation occurs on files that are in use while the OS is running. Files that can be defragmented only while the OS isn't running include the Master File Table (MFT), the hibernation file, and the paging file. Only Diskeeper and PerfectDisk let you schedule an offline defrag.

The products differ in their online defrag approach—that is, defragmenting files while the system is running. Both PerfectDisk and O&O Defrag take a scheduled defrag approach, and they both have wizards to help you automate the scheduling process. Diskeeper constantly runs in the background.

If you need to deploy and manage defrags on multiple systems from one central loca-

SUMMARY

Diskeeper 2008 Server

PROS: System runs in the background and keeps your disks defragged; no need for scheduled defrags; no performance hit

CONS: The most expensive of the three evaluated products; cluttered interface; leaves files fragmented if little free space available

RATING: ◆◆◆◆◇

PRICE: $299.95 per server; volume discounts available

RECOMMENDATION: Install it, and have a nice day. No further action is required. Now *that's* commendable.

CONTACT: Diskeeper • www.diskeeper.com • 818-771-1600

Table 1: Hard Drive Defrag Metrics

	BEFORE	Diskeeper	O&O Defrag	PerfectDisk
Time to Manual Defrag		2 minutes	4 minutes	18 minutes
Total Fragmentation	23%	23%	23%	23%
File Fragmentation	47%	46%	46%	46%
Total Fragmented Files	19	1	1	1
Total Excess Fragments	476	9	24	2

tion, take a close look at each product's functionality in this area, because each one offers something different. O&O Defrag includes a Network Management tool (with the Server edition), Diskeeper offers an add-on product called Diskeeper Administrator, and PerfectDisk (as of this writing) is working on a new product called Command Center. (According to the company Web site, Command Center will be available to PerfectDisk 2008 customers at no additional charge.)

PerfectDisk and Diskeeper both work with Microsoft Volume Shadow Copy Service (VSS). When a disk defragmenter moves enough data around the disk, VSS can mistakenly assume that a file change has occurred, and thus take a snapshot. Both of these products let you to make the system "VSS aware" so that unnecessary snapshots don't occur.

Diskeeper

I started the application by double-clicking the desktop icon. The opening screen is immediately overwhelming, particularly compared with the GUIs of the other two products. Instead of easy-to-read labels, Diskeeper has cryptic icons that you must "hover" your mouse over to see what they do. Perhaps I'm being picky, but my initial impression was that the other two products offer much simpler interfaces. On a brighter note, a Quick Start Guide in the interface's left pane helps get you started in the right direction. Figure 1 shows the Diskeeper interface.

Defragmenting. Diskeeper is unique in that it offers not only classic online and offline defragmentation that you can set manually and schedule, but also a new method called Automatic Defragmentation (which debuted in Diskeeper 2007). Automatic Defragmentation runs silently in the background to ensure that all your hard disks stay defragmented. My first concern was that this feature would consume valuable resources from the server. But Diskeeper uses InvisiTasking technology to monitor disk I/O, memory allocation, and CPU usage to ensure that Diskeeper never negatively affects your users. Automatic Defragmentation will even choose the appropriate engine to use, depending on the kind of fragmentation you have (e.g., heavy fragmentation, low free disk space).

Letting disk-defragmenter software run in the background and configure itself is a new concept to most administrators. Offline and online manual defragmentation is available, but it isn't necessary to run; you can just install it and have a nice day.

Above and beyond. What sets Diskeeper apart from the other two products is the Automatic Defragmentation feature, ensuring that your systems are always in an un-fragmented state. Another interesting feature is Intelligent File Access Acceleration Sequencing Technology (I-FAAST). This feature, according to Diskeeper, sequences files to take best advantage of the logical and physical characteristics of a volume. In short, Diskeeper orders data on the disk so that the content you use most often can be retrieved faster.

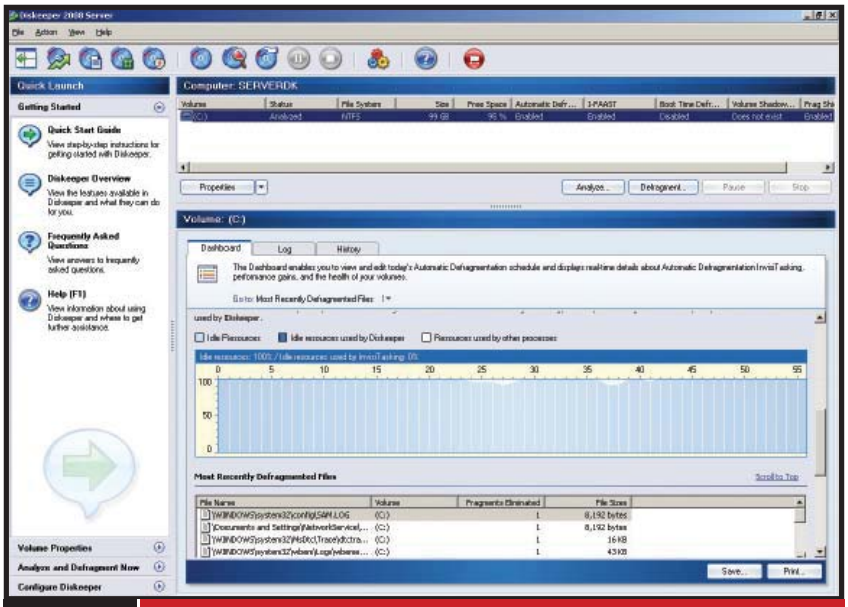

Figure 1: Diskeeper's interface

O&O Defrag

The installation of O&O Defrag proceeded without a hitch. One interesting feature of the installation routine is its *Register O&O Defrag as the standard defrag tool* check box. By contrast, PerfectDisk doesn't replace the default, built-in defragmentation tool that comes with Windows, and Diskeeper replaces it without asking. O&O Defrag gives you a choice.

After installation, O&O Defrag immediately started a wizard to help set up OneButtonDefrag (which Figure 2 shows), a feature that promised to "automate defragmentation with just a few mouse clicks." Opening the other products, I felt unsure where to begin;

Figure 2: O&O Defrag's OneButtonDefrag

O&O Defrag got me started quickly on the right foot. I chose to use the wizard, selected Server (from a choice between Desktop, Laptop, or Server), and then File Server (from a choice between File Server, Database Server, Web Server, or Mail Server).

Next, I used the *Check for Updates* tool in the Help ribbon. The version I'd installed was the most up-to-date version, but this was a great feature that the competing vendors in this category would do well to consider.

Defragmenting. For online defrags, O&O Defrag uses a technology called ActivityGuard that monitors your CPU usage. When you're performing CPU-intensive activities during a scheduled defrag, O&O Defrag uses less of your CPU. But when the processor is idle, O&O Defrag takes the opportunity to optimize more files. You can further tone ActivityGuard to use all available resources or a percentage of the CPU that you specify. In addition, you can set O&O Defrag to defrag each physical drive simultaneously or sequentially. You might choose simultaneous defrags if you need to quickly defrag all of a server's disks. Defragging sequentially takes more time but saves valuable system resources.

Offline defragmentation occurs at boot time. It's disabled by default, so you have to specifically set it by clicking Settings and accessing the Offline Defragmentation tab. You can set the tool to defrag at every startup or on just the next startup. The defragmentation occurs right after a Chkdsk.

OneButtonDefrag is a great way to quickly set up online defragmentation. O&O Defrag sets up the schedule and all the options for you. But if you want more control of how your system optimizes its files, you can set everything manually. Adding your own job in this way reveals many of the advanced features that OneButtonDefrag takes care of for you. There are five available defragmentation methods that you can choose from (i.e., Stealth, Space, Complete-Access, Complete-Modified, Complete-Name), depending on the available resources of the server, the amount of files and free space on the hard disk, and the system's primary use. For example, the Complete-Access method places recently used files at the beginning of the partition, thereby reducing access time.

The two remaining tabs are for scheduling scripts to run either before or after a scheduled defrag. According to the user manual, the scripting feature can be useful for shutting down applications such as Exchange Server or SQL Server before a defrag run, then starting these services back up again.

Above and beyond. O&O Defrag has the simplest interface of all three products (the PerfectDisk coming in a close second). The OneButtonDefrag wizard helps you ensure that you set up your defrag schedule correctly the first time.

If I have one complaint, it's that O&O's support is lacking. The only number on the O&O Web site is German, and I couldn't get through after repeated attempts. The support Web site doesn't offer much information, either. For example, O&O Defrag has provisions to run CMD scripts before and after a scheduled defrag. I'm familiar with writing CMD scripts to shut down and restart NT services, but some administrators might not know where to start. The addition of a knowledge base to discuss this kind of problem would be a great benefit.

PerfectDisk

Like the other two products, PerfectDisk boasts a simple setup routine, asking basic questions and proceeding smoothly. The PerfectDisk installer comes in an MSI format suitable for deployment via your favorite method (e.g., Group Policy, Microsoft Systems Management Server—SMS). PerfectDisk is also written to be controlled through a Group Policy Administrative Template (ADM). So, not only can you deploy the application to your other servers and workstations, but you can control what those users can do with PerfectDisk.

I started the application by double-clicking the desktop icon. Doing so brought up the main PerfectDisk window.

Defragmenting. When I first started PerfectDisk, I needed a little direction. I perused the user guide on the CD-ROM and checked out the company Web site, but I got better information when I contacted tech support. A friendly technician directed me to a knowledge base article titled "How Often

SUMMARY

O&O Defrag 10 Server

PROS: Built-in network-management console; OneButtonDefrag; AutoUpdate feature ensures that you always have the latest version

CONS: Poor support page; no toll-free support number; had difficulty defragging large files with little disk space left

RATING: ◆◆◆◇◇

PRICE: $249 per server; volume discounts available

RECOMMENDATION: If you need network defrag management but don't want to pay extra for it, O&O Defrag gets my recommendation—but if you live in North America, you might have support problems.

CONTACT: O&O Software • www.oo-software.com • (49) (30) 4303-4303

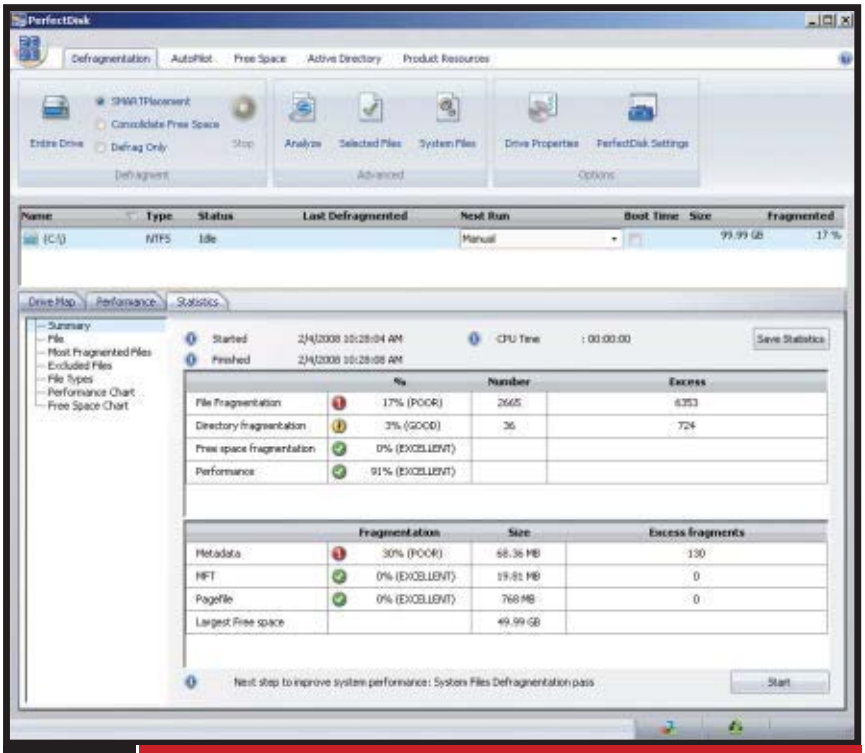

Figure 3: PerfectDisk's interface

Should I Defragment My System?" This brief article suggested performing a drive analysis to see what kind of defrag PerfectDisk recommends.

This analysis took only a few minutes, and at the end, a Start button appeared in the screen's lower right corner, as you see in Figure 3. Clicking this button brought up a cryptic dialog box that read, *Offline defrag of your System Files could not run on drive C: because the drive is in use by another process. Do you want to force all open handles closed?* Like the other products in this review, PerfectDisk can't defrag system files such as the paging file and MFT because they're in use. I expected PerfectDisk to ask me whether I wanted to schedule a defrag at system reboot, but strangely, PerfectDisk attempted to close those system files, then prompted me to reboot the system so that the offline defrag could proceed.

After the reboot, I returned to the Analyze screen and PerfectDisk prompted me to analyze the disk again. I did so, then clicked the Start button, and an online defrag started immediately. According to the technician that I talked to on the phone, running an offline defrag followed by an online defrag is the recommended approach to drive optimization. After these two processes run, you simply need to schedule an online defrag. You can also set up a manual schedule. You can choose the drives to include in the schedule, the defragmentation type, and the date/time you want the defrag to run.

Above and beyond. I appreciated PerfectDisk's ability to schedule an offline defrag (it will automatically reboot the server for you). The product can also "pause" the offline defrag after it has finished so that you can see the results. Active Directory (AD) integration lets you not only deploy the software but also configure it through Group Policy.

SUMMARY

PerfectDisk 2008 Server

EDITORS' CHOICE Windows IT Pro

PROS: The most inexpensive of the three evaluated products; AD integration and deployment

CONS: Unintuitive console has a tough time with disks that have limited free space

RATING: ◆◆◆◆◆

PRICE: $99 per server; volume discounts available

RECOMMENDATION: PerfectDisk is an outstanding value, earning my highest recommendation and *Windows IT Pro*'s Editor's Choice distinction.

CONTACT: Raxco Software • www.raxco.com • 800-546-9728

Results

All three of these disk-defrag products installed flawlessly and worked as advertised. Each defragmented the very full test hard disk completely, except for a large 3.5GB file with 19 fragments. All the products struggled with that file, partially because there was little free space to work with. PerfectDisk struggled a bit more than the others on disks with little free space. The products did equally well on the disk with lots of free space, each removing all defragmentation in about 20 minutes.

I was impressed by Diskeeper's innovative approach to keeping the hard disk constantly defragged. Diskeeper claims that although you can manually defrag and even run the tool under a set schedule, it isn't necessary because the application constantly defrags in the background. Diskeeper cruised through the manual defrag but left half the fragments of the large file. The product's high sticker price, along with the equally high price of the add-on Administrator tool, keeps Diskeeper from attaining the Editor's Choice distinction.

O&O Defrag's poor online and over-the-phone support infrastructure damages the tool's overall effectiveness. The company needs to implement a phone number that's easier to call from North America, and it needs to enhance its support Web site. O&O Defrag also had trouble defragmenting the large file, actually raising the number of fragments.

Unlike Diskeeper's continuous defragging, PerfectDisk uses a manual/scheduled defrag routine that's similar to that of O&O Defrag, so you'll have to schedule defrags. Although it's the slowest of the three products, PerfectDisk does the best job of defragging the hard disk. The extremely low price—combined with its full feature set and free, forthcoming Command Center—sets it apart from the competition and earns PerfectDisk our Editor's Choice award.

InstantDoc ID 98577

Eric B. Rux

(ebrux@whshelp.com) is cofounder of WHSHelp.com. You can read his monthly column "Coming Home to Windows Home Server" at www.connectedhomemag.com. Eric is a senior Windows administrator and teaches the Microsoft Certified Systems Administrator (MCSA) program at a tech college.

Enterprise Antivirus Software

Protect your network

I would venture to guess that virtually every computer network has had to deal with the downtime and expense of recovering from some type of malware infection. According to AV-Test (www.av-test.org), an independent antivirus software testing lab, 2007 saw record numbers of computer viruses, worms, and other malware, and 2008 is continuing that trend. Naturally, prevention is less costly than recovery—but how do you choose from the myriad of antivirus or anti-malware solutions on the market? Let's look at some things you should consider when choosing an enterprise antivirus product, and then you can check out the product comparison table to find the best one for your organization.

Choices, Choices

Today's antivirus market includes products that protect file servers, email gateways, Web browsers, and desktops. They may be standalone products or part of an integrated security suite that might include a firewall, intrusion detection system (IDS), intrusion prevention system (IPS), Network Access Control (NAC), and spam filtering. You can choose from desktop solutions or server-side solutions that offer centralized control for deploying, configuring, and updating the software and that eradicate malware threats before they infiltrate your network. Security appliances as well as hosted and managed security solutions that outsource the management details of your security strategy are also gaining in popularity. Because of the wide array of solution types, we've limited the scope of this Buyer's Guide to server-side enterprise antivirus products.

Features and Functionality

At a minimum, your antivirus solution needs to be compatible with your enterprise OSs and be able to scale and grow with your organization's needs. It should provide frequent automatic signature updates and alert generation when an event is detected. In addition to detection, your solution should provide quarantine or removal functionality and perhaps healing capabilities for suspicious content. Antivirus technology is continuously evolving, so here are some additional features and functionality you should keep in mind.

Scanning engines—the more the merrier. Many antivirus solutions use more than one engine to scan for security threats. No antivirus scanning engine catches 100 percent of viruses. Therefore, using a product with multiple scanning engines can usually pick up the occasional virus or worm that might sneak by a single-engine product.

Detection types—keeping up with new viruses and variants. Most antivirus products detect viruses by using signature-matching technology, which identifies a virus by a specific code sequence. But in today's fast-evolving security environment, when new virus variants crop up by the minute, signature matching isn't enough. Many products now use heuristic scanning and behavior monitoring to identify typical infection methods and suspicious behavior that might indicate virus variants before a signature is available. Unfortunately, these methods can also provide a high number of false positives.

Scanning options—what, where, when. Antivirus products should scan memory, all drives, and the registry. Many now offer scanning of removable devices such as USB drives. They should offer scheduled scans and on-demand scans, and many offer continuous background scanning. Another useful feature is the ability to whitelist items to be ignored or excluded during scans. Reports of the scan log files should be available or portable to your desired format. Reports are important tools for letting you see how many and which viruses have been blocked and where the most popular sources of infection are.

Viruses, worms, and Trojans, oh my. Simply detecting and blocking a virus in an email is no longer sufficient. An antivirus program should detect viruses, worms, Trojan horses, Web threats, rootkits, and other forms of malware that threaten your network security. Your solution should also give you the ability to block certain file types such as .exe, .bat, or .asp files.

Do the Legwork

Of course the most important evaluation criterion for an antivirus solution is performance: high threat detection rates, with few false positives and low impact on business operations. However, performance is beyond the scope of this Buyer's Guide, so we'll leave that part of the evaluation to you. But fear not, there's help. Antivirus testing labs such as AV-Test, ICSA Labs (www.icsalabs.com), and AV-Comparatives.org (www.av-comparatives.org) have done the performance testing for you. So after you have your short list of products that best meet the needs and wants of your organization, visit one or more of these sites for help in determining how the products stack up against one another performance-wise. And don't forget, most vendors (including all those listed in the product table) offer fully functional trial versions so you can try before you buy.

InstantDoc ID 98441

Gayle Rodcay (gayle@windowsitpro.com) is a senior editor for *Windows IT Pro* and *SQL Server Magazine*, specializing in Microsoft Office and SharePoint. She has worked as a technical editor since 1992.

Company	Product	Price	Windows OSs	64-bit Support?	Detection Type	Malware Detected	Scanning Type (Continuous, On-Demand, Scheduled)
AEC 420-541-235-466 www.aec.cz/index.php?english	TrustPort Workstation	$39 for 1 license, $89 for 3 licenses, $143 for 5 licenses	Windows Vista, 2003, XP, 2000	No	Pattern matching, heuristic analysis	All types	C, O, S
AVG Technologies 321-274-1888 www.avg.com	AVG Internet Security 8.0	$54.99 for 1 year; $32.99 renewal	Windows Vista, Vista x64 Edition, XP, XP Pro x64 Edition, 2000	Yes	Pattern matching, heuristic analysis	All types	C, O, S
eEye Digital Security 949-900-4100 www.eEye.com	Blink Enterprise Edition	$69	Windows 2008, Vista, 2003, XP, 2000, NT 4.0	Yes	Pattern matching, heuristic analysis, sandbox, IPS	All types	C, O, S
ESET 866-343-ESET 619-876-5400 www.eset.com	ESET NOD32 Antivirus Business Edition	$32.99 for 1 year; $22.99 renewal	Windows Vista (32-bit and 64-bit), XP (32-bit and 64-bit), 2000	Yes	Pattern matching, heuristic analysis	All types	C, O, S
	ESET NOD32 Antivirus for Microsoft Windows File Server	$20 for 1 year; $14 renewal	Windows 2003, 2000	Yes	Pattern matching, heuristic analysis	All types	C, O, S
F-Secure 888-432-8233 408-938-6700 www.f-secure.com	F-Secure Anti-Virus for Windows Servers	$148	Windows Vista, XP, 2000	Q4 2008	Pattern matching, heuristic analysis, behavioral-based	All types	C, O, S
Kaspersky Lab 866-328-5700 781-503-1800 www.kaspersky.com	Kaspersky Anti-Virus 6.0 for Windows Servers Enterprise Edition	Contact vendor	Windows 2003, 2000, and associated terminal servers	Yes	Pattern matching, heuristic analysis	All types	C, O, S
Microsoft 800-642-7676 www.microsoft.com	Forefront Client Security	Contact vendor	Server: Windows 2008, 2003 Client: Windows Vista Business, XP SP2, 2000 SP4	Yes	Pattern matching, heuristic analysis, tunneling signatures for detecting and removing rootkits	All types	C, O, S
PCSecurityShield 866-615-7443 561-243-3247 www.pcsecurityshield.com	The Shield Deluxe 2008	$39.99	Windows Vista, XP, 2000	Yes	Pattern matching, heuristic analysis	All types	C, O, S
Sunbelt Software 888-688-8457 727-562-0101 www.sunbeltsoftware.com	CounterSpy Enterprise	$29 per seat (up to 25 seats); sliding scale discount	Windows Vista, XP, 2000 SP4 and later	Yes	Pattern matching, heuristic analysis, function matching, behavior	All types	C, O, S
Symantec 800-745-6054 408-517-8000 www.symantec.com	Symantec Endpoint Protection 11.0	$51.60 for 5-24 licenses, $41.60 for 24-49 licenses, $39.75 for 50-99 licenses, $37.90 for 100-249 licenses, $35.00 for 250-499 licenses, $31.80 for 500-1,000 licenses	Windows Vista, 2003, XP, 2000	Yes	Pattern matching, heuristic analysis	All types	C, O, S
Trend Micro 877-218-7363 408-257-1500 www.trendmicro.com	InterScan Messaging Security Suite	$7.96-$42.77	Windows 2003 R2, 2003, 2000	Yes	Pattern matching, heuristic analysis	All types	C, O, S
	ScanMail for Microsoft Exchange	$12.50-$43.47	Windows 2003 R2, 2003, 2000	Yes	Pattern matching, heuristic analysis	All except keyloggers	C, O, S
	InterScan Web Security Suite	$10,560 (1,000 users)	Windows 2003, 2000	No	Pattern matching, heuristic analysis, real-time URL reputation	All types	C
	OfficeScan 8.0	Contact vendor	Windows Vista, 2003, XP, 2000; Windows Cluster Server 2003, 2000; Windows Storage Server 2003	Yes	Pattern matching, heuristic analysis (IntelliTrap)	All types	C, O, S
Webroot Software 800-870-8102 303-442-3813 www.webroot.com	Webroot AntiSpyware Corporate Edition with AntiVirus	$39.95	Windows Vista, 2003, XP, 2000; VMware Workstation 5.5 or later	No	Pattern matching, heuristic analysis, sandbox, direct disk access for rootkit detection	All types	C, O, S

EDITOR'S NOTE: Some vendors that you might expect to see in this Buyer's Guide said they didn't have a product that exactly matched the criteria

	Signature Updates (Hourly, Daily, On-Demand)	Scanning Engines	Scan Location (Memory, Registry, All Drives, Removable Media)	Report Options	Automatic Alerts?	Remote Management?	Supports n-Tiered Deployments?	Centralized Management, Monitoring, and Reporting?	Fully Functional Trial Version?
	H, D, O	Norman, Dr Web, AVG, Ewido, VirusBlockAda	M, R, D, RM	HTML	Yes	Yes	Yes	Yes	Yes
	H, D, O, scheduled	AVG, XPL	M, R, D, RM	Text for integration into third-party reporting packages	Yes	Yes	No	Yes	Yes
	H, D, O	Signatures, IPS, sandbox, eEye heuristics, phishing	M, R, D, RM	PDF, Excel, Word, HTML, ActiveReports, CSV, XML	Yes	Yes	Yes	Yes	Yes
	H, D, O, as needed	ESET NOD32 Antivirus	M, R, D, RM	Copy and paste from application	Yes	Yes	Yes	Yes	Yes
	H, D, O	ESET NOD32 Antivirus	M, R, D, RM	Copy and paste	Yes	Yes	Yes	Yes	Yes
	H, D, O, automatically	Orion, Libra, AVP, Blacklight, Gemini, Pegasus, Draco	M, R, D, RM	HTML, log files, system event files	Yes	Yes	Info not provided	Yes	Yes
	H, D, O, automatically	Kaspersky Anti-Virus engine	M, R, D, RM	Excel, HTML, email, NETSEND, run script	Yes	Yes	Yes	Yes	Yes
	H, D, O	Microsoft Malware Protection Engine	M, R, D	PDF, HTML	Yes	Yes	Yes	Yes	Yes
	H, D, O	Kaspersky Anti-Virus engine	M, R, D, RM	Text	Yes	No	No	No	Yes
	H, D, O	VIPRE hybrid antivirus/ anti-malware scanning engine	M, R, D, RM	PDF, Excel, HTML, RTF, TIFF, TXT, custom	Yes	Yes	No	Yes	Yes
	H, D, O	Info not provided	M, R, D, RM	HTML, text, exported SHTML	Yes	Yes	Yes	Yes	Yes
	H, D, O, weekly, by minute	Spam engine, malware engine, spyware engine	N/A	HTML, CSV	Yes	Yes	Yes	Yes	Yes
	H, D, O, by minute	Spam engine, malware engine, spyware engine	N/A	HTML	Yes	Yes	Yes	Yes	Yes
	H, D, O, every 15 minutes	All Trend-developed antivirus, antispyware, antiphishing, anti-malware, and URL categorization engines	N/A	HTML, CSV	Yes (email and SNMP)	Yes (HTTPS)	Yes	Yes	Yes
	H, D, O	VSAPI virus engine, SSAPI spyware engine, RCM rootkit engine, DCE virus cleanup engine, TMUFE URL filtering engine, IntelliTrap engine	M, R, D, RM	PDF, Excel, Word, HTML	Yes	Yes	Yes	Yes	Yes
	D, O	Sophos, Webroot	M, R, D, RM	PDF, Excel, RTF, CSV	Yes	Yes	Yes	Yes	Yes

or didn't respond to our requests for information about their products.

MASTERING RSoP

Your first step to Group Policy health

by Darren Mar-Elia

I spend a fair bit of time helping folks figure out problems with Group Policy. In the 8+ years I've been doing this, by far the biggest improvement in Group Policy management has been the introduction of the Resultant Set of Policies (RSoP) capability in Windows Server 2003 and Windows XP to help us figure out what the effective policy is on our desktops and servers. Understanding what RSoP is and knowing how to read an RSoP for a user or computer will help you ensure Group Policy is healthy and happy in your environment. And, while RSoP won't help solve every Group Policy problem that arises, an RSoP can point the way toward how to further investigate.

What Is RSoP?

The first thing to understand about the RSoP feature in Windows is that it's technology that Microsoft built into the Windows Management Instrumentation (WMI) infrastructure beginning with Windows XP. RSoP doesn't support Windows 2000 because Win2K's WMI infrastructure and Group Policy engine don't include the necessary components to collect RSoP information. The Windows 2000 Resource Kit does ship with a command-line utility called gpresult.exe that provides some of the information that RSoP delivers, but this first-try Gpresult doesn't paint as complete a picture of policy processing as the later RSoP.

When XP and later versions of Windows were introduced, Microsoft provided two main tools for accessing the WMI-based RSoP infrastructure. The Microsoft Management Console (MMC) Group Policy Management Console (GPMC) snap-in provides a graphical UI for accessing RSoP data, and the command line–based Gpresult is built into the OS. Don't confuse the RSoP-enabled version of Gpresult with the earlier Win2K Resource Kit version. Because the two tools use completely different mechanisms, they can return different information, with the RSoP-enabled version being the more accurate of the two.

So what exactly is RSoP? Well, essentially it's a mechanism to determine, for a given computer or user in Active Directory (AD), what that computer or user's effective Group Policy settings are. A user or computer can process many Group Policy Objects (GPOs) in a typical AD environment—with GPOs having possibly conflicting settings. GPOs are processed in a certain order that affects which policy settings will actually apply to a given user or computer, and GPOs can be filtered by using security groups and WMI filters. Given all these factors, you can see how knowing what the effective policy settings are for a given user or computer can be hard, especially in larger organizations. RSoP cuts through the confusion and tells you what's happening with your Group Policy settings.

RSoP Planning vs. Logging

The RSoP capability in Windows Server 2008, Windows Vista, Windows Server 2003, and XP comes in two flavors. The first, and by far the most common, is known as RSoP or Group Policy Results Logging. (Group Policy Results is the more common name for RSoP.) Group Policy Results Logging, as the name implies, lets you see what policies were applied to a given Windows computer or user. It answers the question, "What policy settings were processed by a given computer or user during the last policy processing cycle?" Logging relies on the Group Policy engine and each Client Side Extension (CSE) that processes the various policy settings to report to WMI on what it did when Group Policy was processed. When you run a GPMC Group Policy Results Logging report, which Figure 1, page 34, shows, or use Gpresult from your XP or Vista machine, you're essentially connecting to the machine you select—local or remote—and gathering the WMI logging data into a report.

The second RSoP flavor, RSoP Planning (also known as Group Policy Modeling in GPMC), answers the question, "What policy should apply to a given computer or user during a future policy processing cycle?" As the name implies, RSoP Planning lets you perform a "what-if" calculation on the policy that a given

Learning Path

WINDOWS IT PRO RESOURCES

For help troubleshooting Group Policy:

"Troubleshooting Group Policy–Related Problems," InstantDoc ID 44983

To use RSoP from the command line and in scripts:

"How can I view the Resultant Set of Policy (RSOP) from the command line?" InstantDoc ID 97129

"GPMC Scripting," InstantDoc ID 39529

To gather Group Policy deployment info for multiple computers or users:

"Monitor GPO Deployment," InstantDoc ID 50256

computer or user will receive. It goes one step better and lets you play with changes that might occur to users or computers to see what effect the changes will have on the users' or computers' effective policy.

For example, you can virtually move the user or computer into a new organizational unit (OU) or new security group to see how that will impact its effective policy. You can also simulate how policy would be affected if a slow network link were detected or if loopback policy were in place. All of these "modifications" that you can perform during the modeling phase will affect what policy settings a computer or user receives, and the Group Policy Modeling feature in GPMC lets you simulate these changes easily.

Unlike Group Policy Results, Group Policy Modeling doesn't require you to query a particular target computer to figure out what will happen. However, it does require access to a Server 2003 or Server 2008 domain controller (DC) to work. In fact, if you have only Win2K DCs in your AD domain, you won't even see the Group Policy Modeling node when you start up GPMC because the modeling feature uses a service called the Resultant Set of Policy Provider that runs only on the newer DCs. Without this service, modeling won't run.

Using and Deciphering RSoP Logging

Now that you know what RSoP is, let's look at how you can use it to get more insight into your Group Policy settings. I find the version of Group Policy Results Logging that's available in GPMC easier to use than the command-line Gpresult utility, so let's start with the graphical version.

A note before we dive into the details: If you're working in a mixed environment of Server 2008, Vista, Server 2003, and/or XP, the rule of thumb is to run Group Policy Results on a machine that has the same or a more recent OS version than the machine whose results you're testing. So, if you're running Group Policy Results against a Vista machine, run it from a Vista machine, not an XP machine. You'll get more complete results this way.

The other thing to be aware of at this point is that the computer for which you're collecting Group Policy Results must be accessible on the network from your management station. That means it must be up and running and must not have a firewall blocking access to the ports and protocols required by Group Policy Results. Because Group Policy Results uses remote WMI calls to get access to this information, you typically need to ensure that the remote system allows the DCOM protocol. This protocol uses TCP port 135 for initial setup and random ports greater than 1024 for ongoing communication. If the target machine uses Windows Firewall, the easiest way to ensure that the necessary ports are unblocked is to use the built-in Remote Administration Exception provided in Group Policy. You can find this exception on XP and Server 2003 in GPMC under Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall\Standard (or Domain) Profile\Windows Firewall: Allow Remote Administration Exception and on Vista and Server 2008 under Computer Configu-

IT Pro Hero

Keys to Group Policy Success? Prepare and Test!

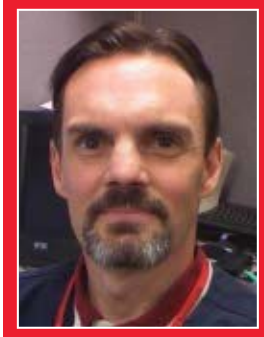

LAN administrator Mike Foster gives advice on how to succeed using GPOs to deploy software

BY CAROLINE MARWITZ

Group Policy users typically can tell at least one horror story about settings gone wrong, but 10-year IT veteran Mike Foster says he hasn't really had what he'd call a horror story happen, even when his organization used Group Policy to install Sun Microsystems' Java Runtime Environment (JRE) on 800 computers. The success of that experience was due in no small part to Mike's preparatory steps. Mike is currently a junior LAN administrator for a US government organization that focuses on health care, but he gained his background in Group Policy as a Microsoft Certified Trainer working with Active Directory (AD). Besides asking him to share his experience with using Group Policy to install JRE, *Windows IT Pro* Web Site Strategic Editor Anne Grubb and I quizzed Mike about how to get up to speed with Group Policy resources. We even managed to glean an almost-horror-story from him, which he diplomatically calls a Group Policy "challenge," about deploying software at remote sites.

Q: Your organization needed to install JRE on 800-plus computers. How did you use Group Policy in this situation?

A: Deploying JRE was fairly straightforward. For me, the biggest hurdle was extracting the .msi file from the JRE installation executable (.exe) file. The .msi file is required to do a Group Policy–based installation. For help with extracting the .msi file and other aspects of the JRE deployment, I referred to Sun's documentation at java.sun.com/j2se/1.5.0/docs/guide/deployment/deployment-guide/upgrade-guide/deployment.html and java.com/en/download/help/5000011100.xml.

Once I obtained the required .msi file, I simply followed best practices by assigning the application in the Group Policy Computer Configuration/Software Settings/Software Installation node in Group Policy Editor and verifying

that I placed the required .msi file in an appropriately shared folder in an accessible network location, with the correct permissions configured on it. We didn't require any scripts or transform file for this installation, but we've used user logon scripts with Group Policy to configure the user environment, such as for modifying registry values required by applications or the users. We've also used computer startup scripts with Group Policy for various purposes.

Q: Was there anything special you did on this project that helped to make things work smoothly?

A: I conducted thorough testing on each client platform prior to rolling this out to our production environment. I also notified all users about the rollout in advance because Group Policy software installations occur during the computer boot phase, which leads to a delay when booting the computer. We notified our users well in advance so that calls to our Help desk would be minimized during the installation phase.

Q: What advice would you give others looking to deploy applications by using Group Policy?

A: First I recommend conducting thorough research in advance, so that you completely understand what the requirements are before you get started. Review white papers and best practices for using Group Policy. These are available via Microsoft's Web site and elsewhere on the Internet. I also recommend thoroughly testing the policies and deployment in a test environment on each client OS used in your production environment to ensure there are no issues once you get to production.

Q: Do you have any Group Policy horror stories?

A: I wouldn't call it a "horror story," but for me one of the biggest challenges with using Group Policy to perform software installations has been our remote sites and the WAN bandwidth issues we face. I created and targeted specific organizational units (OUs) within our AD so the computers at the remote sites wouldn't get large software installations over the WAN. We did do some smaller installs over the WAN, such as our Daylight Saving Time patch, but for the most part I recommend using multiple Group Policy Objects (GPOs), each with its own localized software source directory, and targeting specific OUs based on geographic location. Or you could do large software deployments manually on each client at the remote sites (which is something that we did for small remote sites where the bandwidth didn't support a software rollout using Group Policy.)

Q: What do you think of the Resultant Set of Policy (RSoP) snap-in?

A: I've used RSoP to troubleshoot Group Policy configuration settings, and I've also used the GPResult command-line tool. As somebody who came into the IT arena after the invention of the GUI, I really appreciate tools such as RSoP because I tend to grasp the information quicker from a GUI as opposed to the results of command-line tools such as GPResult. One of the benefits I've seen with RSoP is that it allows for reviewing the existing GPOs that are applied to a given computer and/or user (logging mode), which is great when you're troubleshooting Group Policy settings. RSoP also provides a way for the administrator to simulate the effect of applying a GPO (using planning mode), without actually applying the policy to the target computer and/or user.

InstantDoc ID 98477

MORE ON THE WEB

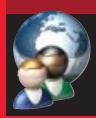

Read an expanded version of this article at www.windowsitpro.com, InstantDoc ID 98477.

Caroline Marwitz
(cmarwitz@windowsitpro.com) is an associate editor for *Windows IT Pro* and *SQL Server Magazine*, specializing in Active Directory, Group Policy, and desktop management.

ration\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Inbound Rules, under the Predefined Rules selection.

Ok, let's get started. Suppose you want to verify that a certain workstation has retrieved some policy settings. Start GPMC, right-click the Group Policy Results node, then select the Group Policy Results Wizard option. The first wizard screen lets you select a remote or local computer to connect to. If you're interested only in per-user settings, you can also select a check box to exclude any per-computer settings in the report that will be generated.

After selecting the computer you want to target, the next wizard screen lets you select a user who has logged onto that computer, if you want to return per-user Group Policy settings in addition to computer settings. The Group Policy Results wizard UI will show you only those users who have logged onto the remote system and generated RSoP data. If you don't see a user in the list, he or she likely hasn't logged onto that system. After you select the user, the Group Policy Results wizard collects the WMI data from the selected computer and displays it in the GPMC's right-hand results pane, as shown in Figure 1.

Interpreting the Results

Once you've run the Group Policy Results wizard and the results are displayed, you can dive in and interpret those results. In the right-hand results pane are three tabs: Summary, Settings, and Policy Events. Table 1, page 36, describes the purpose of each.

The Summary tab is probably the most interesting in terms of finding out what's going on with Group Policy on the remote system, so let's examine it in detail. Figure 2, page 36, shows an expanded Summary tab with all its sections.

Assuming you selected to show both per-computer and per-user Group Policy settings, the summary will be broken into two sections: Computer Configuration Summary and User Configuration Summary. In each of these sections are five subsections that provide details about what policies were processed. The most interesting subsections are Group Policy Objects and Component Status.

The Group Policy Objects subsection is further divided into Applied GPOs and


Your potential. Our passion.™
Microsoft®
MICROSOFT SYSTEM CENTER. DESIGNED FOR BIG.
Microsoft® System Center is a family of IT management solutions (including Operations Manager and Systems Management Server) designed to help you manage your mission-critical enterprise systems and applications.
Dell™ is using System Center solutions to manage 13,000 servers and 100,000 PCs worldwide. That's big. See Dell and other case studies at DesignedForBig.com
Microsoft®
System Center

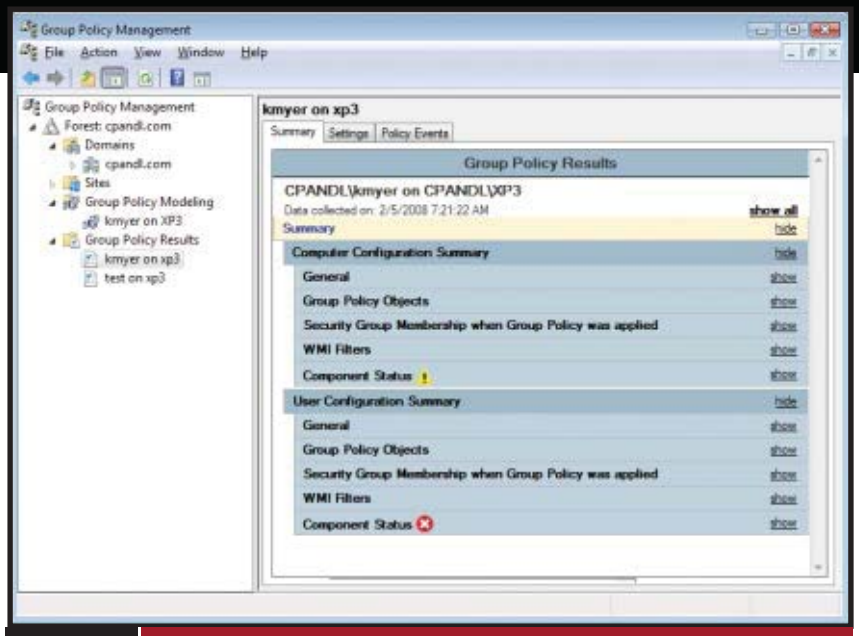

Figure 1: A Group Policy Results report in GPMC

Denied GPOs. Applied GPOs lists the GPOs that were processed by the computer or user, to which AD container those GPOs were linked, and what their AD and SYSVOL version numbers were. This information is important because it lets you verify that a particular GPO that you think should be processed by the computer or user really is being processed. The version numbers are important because they should always be the same for a given GPO. If the AD and SYSVOL version numbers are different from each other, the GPO being processed could be out of sync on the DC that the computer is using to process policy, which could indicate a replication problem (or simply that you initiated Group Policy Results processing without leaving enough time for GPO changes to replicate to the DC).

The Denied GPOs section is equally interesting because it tells you exactly why a GPO wasn't processed, even though it might be linked to a container in AD that includes the computer or user. The most common reasons for GPOs being denied—or, more correctly, not processed—include security group filters or WMI filters that prevent them from being processed, a link being disabled, or the GPO being empty (i.e., containing no settings). The Denied GPOs section can provide good information about how you're applying your policies and might indicate places where you can get rid of "dead" GPOs that computers or users are trying to process but can't.

The Component Status section of the results is the really interesting part! It's the portion of the report that tells you whether Group Policy processing actually worked for the computer or user in question. This section of the report is broken down into each part of the Group Policy processing cycle. The component named Group Policy Infrastructure represents what's called the core phase of Group Policy processing. During this phase, the computer or user account reads the list of GPOs it must process, finds out which ones it has access to, and creates the list of CSEs that must process the policy settings in those GPOs. If this part of the processing cycle fails, then the rest of the cycle will fail.

The subsequent components listed are the various CSEs that ran for the computer or user during the last processing cycle. These include the different policy areas such as Registry, Security, and Software Installation. The report shows whether each item succeeded or failed and, if it failed, will often show the related error information. In Figure 2, the Software Installation item is Pending. The software installation CSE requires a foreground processing event (i.e., a system reboot or user logon) to actually run, so while the component hasn't failed, it hasn't yet run. The Component Status section is also marked with a visual cue—a red X in the case of a failed event or a yellow ! in the case of a warning such as the pending state of software installation.

Table 1: The Group Policy Results Tabs in GPMC

Tab	Purpose
Summary	Provides summary information about which GPOs were and weren't processed by the computer and/or user, and why. Also states whether Group Policy processing actually succeeded on the target system and gives some other summary information, such as the security groups that the computer and/or user belonged to at the time of processing.
Settings	Shows the actual settings applied to the given computer and/or user and the "winning" GPOs that applied them.
Policy Events	Creates a filtered view of the Group Policy–related events from the Application event log on the remote computer.

The Settings tab of the Group Policy Results report, shown in Web Figure 1 (www.windowsitpro.com, InstantDoc ID 98442), gives you a breakdown of the actual policy settings that were applied to the computer or user and the name of the GPOs that delivered those settings. The report in Web Figure 1 shows the details of some Administrative Template Windows Firewall settings that were processed by the client. For Administrative Templates, the report actually includes the Explain text that goes along with the policy to remind you what the policy is for. Note also that, on Vista and Server 2008 systems, the report lets you know that Administrative Template policy descriptions were retrieved from the central store, which is the server-based file-system location where ADMX and ADML files can be kept.

When you select the Policy Events tab of the Group Policy Results report, you see a list of Group Policy–related events that occurred on the remote system. These look and feel like Windows Event Viewer events because that's where they come from. In many cases, the events that are the most interesting are the error or warning events, but frankly, I haven't found much of use in this information, due to the sheer volume of events and the lack of detail about them. However, it's worth looking at this view if you're having problems because some useful information could turn up.

If you want to save the information from the Summary or Settings tabs, right-click over the area of either tab and select Save Report to send the report to an HTML or XML file. The XML file format is useful only

if you plan to repurpose the raw data somewhere else.

You can view a five-minute screencast that shows how to run the Group Policy Results wizard and view the output at wms18.streamhoster.com/pentonmedia/windows/winscreencasts/RSOP-MarElia.wmv.

Under the Covers

GPMC and Gpresult hide the complexity of how RSoP data is collected in WMI, but if you're familiar with WMI and know how to query its contents, you can get directly at the WMI data that underlies those nice RSoP reports. RSoP data is held in a special namespace within WMI specifically for that purpose. Whereas you might be familiar with querying information in the root\CIMv2 namespace, RSoP data is held in root\RSOP. The data is broken down into a number of different classes, each representing different policy areas (e.g., registry, folder redirection, security). Figure 3 shows a view into this namespace through a WMI browser tool called WMIX, which you can download at www.pjtec.com/WMIX.

What you see here is a number of containers in the RSOP namespace. The containers that start with NS followed by a bunch of alphanumeric characters are called RSoP Sessions. They represent me running RSoP reports remotely against this system, called XP3. In Figure 3, I've drilled down into one of these sessions and you can see a number of WMI classes representing the various policy areas that you'd typically look at in an RSoP report. If I viewed the instances on these classes, I would see the raw Group Policy settings data that the GPMC report returned.

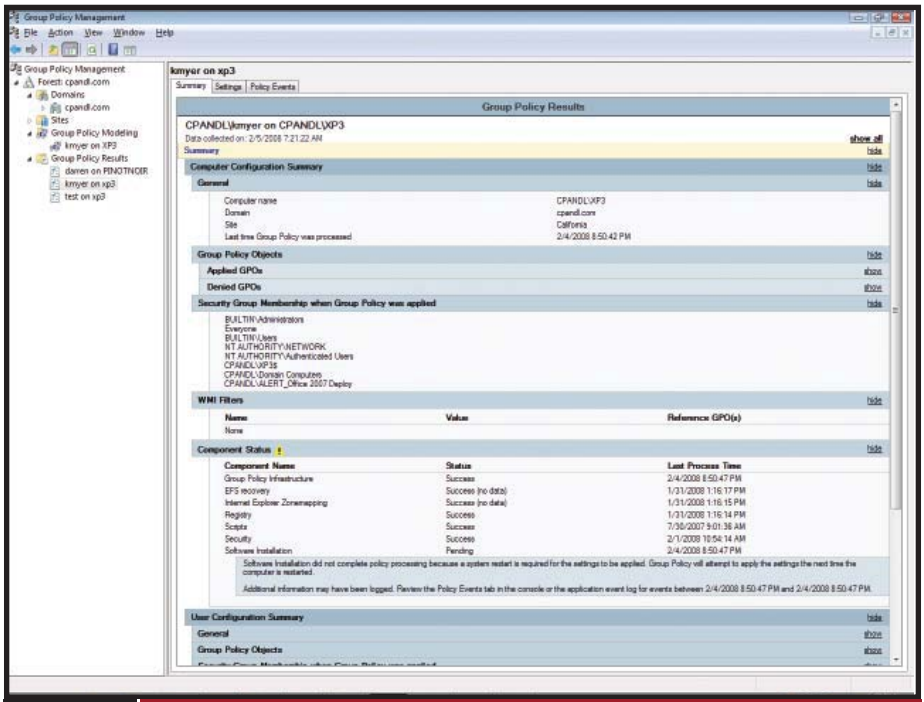

Figure 2: Group Policy Results Summary detail

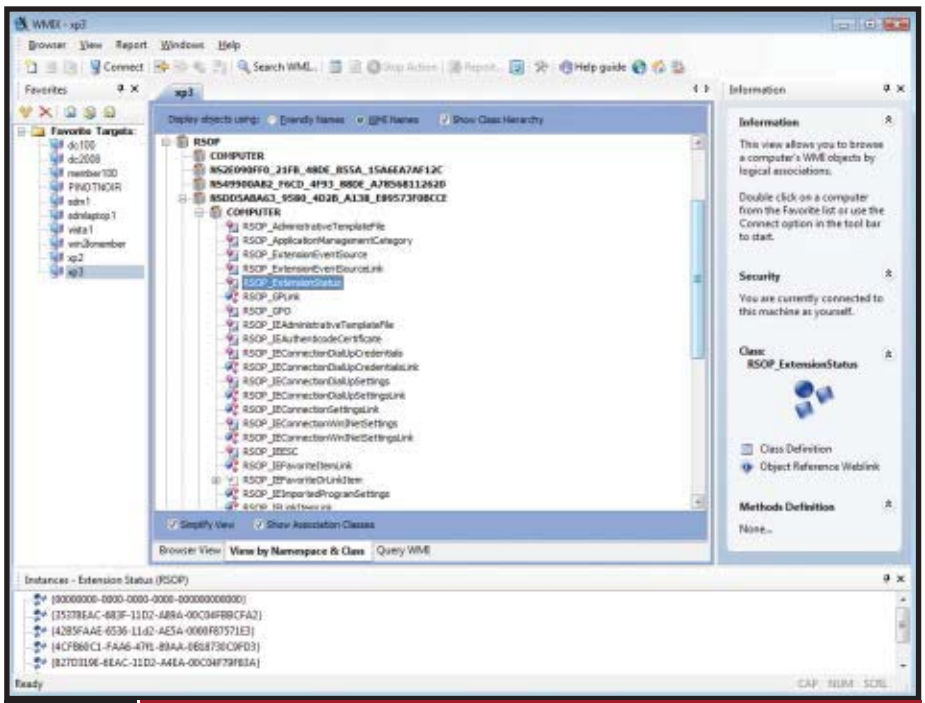

Figure 3: RSoP data in WMI

RSoP data provides a powerful mechanism for discovering how Group Policy is working on your remote Windows systems. Using GPMC or Gpresult, you can both model what should happen with Group Policy for a given user or computer, as well as what did happen. And not only do you get to see the actual settings that were processed, but you can also see whether any errors occurred during processing that might have prevented settings from being delivered. This important tool goes a long way toward guaranteeing that Group Policy is doing what it's supposed to do—keeping your systems secure and locked down.

InstantDoc ID 98442

Darren Mar-Elia

(dmarelia@windowsitpro.com) is a contributing editor for *Windows IT Pro* and is CTO and founder of SDM Software (www.sdmsoftware.com). He maintains a Group Policy resource Web site (www.gpoguy.com) and is coauthor of *Microsoft Windows Group Policy Guide* (Microsoft Press).

EXECUTIVE PERSPECTIVE

AMD Rallies Around Virtualization

Margaret Lewis, Director of Commercial Solutions and Software Strategy, shares AMD's perspective on virtualization today and in the future.

How is AMD positioned in the virtualization market?

Lewis: We've matured the x86 architecture for virtualization since the AMD Opteron processor April 2003 release. Our Direct Connect Architecture, featuring an integrated memory controller, offers the memory access and control needed for memory intensive applications. We've also added hardware capabilities directly in the processor to make virtualization software, like Windows Hyper-V, perform more efficiently.

What is AMD's focus with virtualization?

Lewis: AMD's virtualization design goal is to enable applications to run as efficiently on a virtual machine as they would on a non-virtualized machine. To this end we've added hardware functionality to allow virtualization software to work with less complexity. AMD's virtualization technology, called AMD-V, is the result of this effort, providing software vendors with processor-required functionality to enable efficient and secure virtualization solutions. Windows Hyper-V utilizes these hardware capabilities. We collaborated quite closely with Microsoft as they created Windows Server 2008 with Hyper-V to ensure optimum functionality for our joint customers.

How much will virtualization technologies transform the IT professional's job in the next five to 10 years?

Lewis: Virtualization is enjoying rapid adoption. As IT professionals continue to see the many virtualization benefits, especially in disaster recovery and data management, the technology will become increasingly mainstream. In five to 10 years it will be hard to find a non-virtualized data center. As such, IT professionals will use virtualization to develop more effective data-center resource management and provide new levels of computing resource flexibility. For example, virtualization will make software deployment easier, allow resource shifting to make peak load handling better, and allow more cost-effective management infrastructure.

What impact will multi-core processors have on the growth and development of the virtualization market?

Lewis: Quad-core processors will provide servers more processing capabilities (more density of CPU capabilities per server), which will lead to more virtualization growth and development. This will translate into more efficient virtualized environments. This in turn will help companies consolidate resources, save data-center space, and reduce power and cooling requirements.

What other general technology advancements does AMD see coming for the virtualization market?

Lewis: Server virtualization continues to expand its footprint in enterprise data centers as companies enhance this technology. At AMD, our processors and chip sets have built-in performance accelerator support, such as I/O virtualization and Rapid Virtualization Indexing, so our enterprise customers can leverage virtualization for innovative uses, such as client workloads hosted, in virtual machines. We are also very excited by the potential of hosted-client computing, as it allows the marriage of AMD graphics and low-power client solutions with more server-centric virtualization acceleration, providing very high degrees of flexibility for enterprise computing infrastructure.

Further out, AMD is exploring technologies that will bring virtualization-enabled solutions to mobile computing and desktops. By implementing AMD-V technology for clients, and working on the hard-graphics and I/O virtualization problems, AMD is working to ensure customers can effectively share a media-rich client-computer between multiple virtual machines. Over the next three to five years, customers will see client-based virtualization security, manageability and application-delivery solutions begin to emerge as we bring the silicon feature-set to market and work with virtualization and software vendors, such as Microsoft, to enable this exciting future.

Margaret Lewis *has more than 25 years of experience in developing and marketing commercial computing products. She is currently director of AMD's Commercial Solutions and Software Strategy, where she is responsible for identifying software solutions for AMD's products that target the commercial market. Prior to working at AMD, she was associate director of marketing at the Maui High Performance Computing Center, a government computing center focused on technical computing solutions for the scientific and engineering communities and a marketing manager for Novell's developer and database product groups. Ms. Lewis is a graduate of the University of Texas at Austin.*

System Center to Support Linux and VMware

BY KAREN FORSTER

Long-term success for Microsoft depends on our ability to deliver a platform that is open, flexible, and provides customers and developers with choice. These choices include Microsoft and open source technologies working together ….

—Microsoft's Bill Hilf (general manager, Windows Server Marketing and Platform Strategy) in his February 21, 2008, blog post at port25.technet.com/archive/2008/02/21/interop.aspx.

An exclusive interview with Brad Anderson and Larry Orecklin on announcements at the Microsoft Management Summit

Two highly significant announcements are emerging at this year's Microsoft Management Summit (MMS): In addition to Windows environments, the System Center family will manage various flavors of Linux and UNIX, both physical and virtual; and System Center Virtual Machine Manager (VMM) will manage non-Microsoft hypervisors such as VMware's. These and other major developments demonstrate the company's determination to acquire a competitive advantage by adopting the openness mandated in Hilf's statement.

In an exclusive interview, Microsoft's Brad Anderson (general manager, Windows and Enterprise Management Division) and Larry Orecklin (general manager of System Center and virtualization) disclose groundbreaking announcements and discuss how System Center products will embody and advance Microsoft's strategic commitment to support heterogeneous environments.

Manage Non-Windows Devices

Forster: What MMS System Center announcements do you consider important?

Anderson: One will be our support for systems other than Windows. We'll be

releasing into beta a Microsoft solution that allows System Center to manage non-Windows devices. We're coming to market with the infrastructure that allows System Center to manage Linux, Solaris, and other flavors of UNIX. We'll provide that infrastructure for partners to deliver incremental value on top of System Center. We're leveraging the existing back-end systems of [System Center] Operations Manager and [System Center] Configuration Manager [SCCM] and leveraging an open-source agent on the managed device. We're working with the Open Pegasus community to deliver the integration of that agent into our back-end systems. That's where we'll start to expand System Center capabilities to manage past Windows.

Forster: Allowing System Center to manage heterogeneous systems sounds like the old "embrace and extend" philosophy.
Orecklin: I think it reflects a maturing of the industry and of Microsoft. We recognize that there will be heterogeneity in enterprise environments. As we've become more adopted and accepted in large data centers—we have 75 percent of the servers in data centers today—we believe Windows is the best platform to virtualize your data center. We believe System Center is the best management platform. So we want to ensure it's as easy as possible for customers to adopt that as their standard.

Manage Heterogeneous Hypervisors

Forster: Managing competing OSs is a radical move for Microsoft, but the business advantage is clear. How far will this new commitment to heterogeneity extend?
Anderson: There's another flavor of heterogeneity that we'll also be putting into beta for the next release of VMM. That's the ability to manage hypervisors other than Hyper-V. We'll be able to do the same types of activities, tasks, and scenarios if a customer is using a hypervisor from VMware, using ESX, in the same way we manage Hyper-V. In the same way we're going to extend our reach to monitor and manage more than Windows, we're going to do the same with hypervisors. We'll start with VMware and extend that over time to include XenSource. This VMM beta will have all the capabilities of intelligent placement of virtual machines [VMs] on your hardware, plus physical-to-virtual and virtual-to-virtual migration. You'll have all those capabilities independent of whether you're using Microsoft's or VMware's hypervisor.

Forster: When you say virtual-to-virtual migration, will you be able to migrate a VM running on ESX to a VM running on Hyper-V?
Anderson: The answer is absolutely yes.

Forster: When will this version be available?
Orecklin: We're announcing the beta for the next version of VMM at MMS. It will be released to market in Q3, as soon after Hyper-V as possible.

Forster: Why is interoperability so important for System Center?
Orecklin: The key is that the hypervisor will become ubiquitous. We believe management is that key differentiator that will allow customers to take an interesting fad and make it highly leveraged and valued in their environment. That's why we want System Center to be that point from which you manage the environment. You may have deployed ESX for a couple of workloads. We do not want to force you to change that. But as you add workloads, we want to make it as easy and economical as possible to do that on the Windows platform. And if, over time, you want to change and migrate, fantastic.
Anderson: I'd submit that you cannot achieve the benefits promised by virtualization without that strong management solution. It enables you to recognize the savings. Customers want one cohesive, unified way to manage the virtualized environment and the physical. We'll be able to do that in Windows, non-Windows, with Hyper-V, as well as other hypervisors.

Orecklin: It's all about infrastructure and cost. Think about the skill sets of your readers. They only need one tool to get their job done. They don't have to think about scaling up on different kinds of tools and worrying about is this physical or virtual; is it this environment or that environment?

Server Application Virtualization

Forster: Hyper-V is receiving lots of attention, but Microsoft has also recently made announcements about other layers of the virtualization stack. How does System Center relate to your overall virtualization strategy?
Anderson: We can manage all the way from bare metal through the application or services running inside a VM. Other solutions tout the ability to manage your virtual machine environment, but the reason you deploy a VM is because you have a service or application inside it. We have all the models and knowledge about the applications running in the VMs. We understand Exchange's needs and SQL's needs. There are 200 management packs that are available that run in conjunction with Operations Manager. So our vision is to do your VM management, understanding the needs, characteristics, performance, and capacity of the application in the VM.

Forster: Last February, you announced that Microsoft Application Virtualization (formerly Softricity products) would support virtualized applications on Windows Server, in addition to client-side applications.
Anderson: The focus was how we isolate the application from the OS. How do you take that to the enterprise? SoftGrid will also support server applications, so we'll be able to separate the application from the OS on the server. Now you'll have that flexibility to move applications between servers by keeping them separate.

Forster: How does that affect IT?
Anderson: The number of images IT will have to manage will be dramatically lower. The combination of hardware virtualization and application virtualization will mean IT will have a very small number of OS images—literally a handful. And then they'll have a set of images based on VHD [Virtual Hard Disk]. We'll align the formats for hardware virtualization and application virtualization, and you'll have a set of images of just the

Learning Path

WINDOWS IT PRO RESOURCES

"Open Source Support: What's in It for Microsoft?" InstantDoc ID 95352

"Open Source and Windows Server's Direction," InstantDoc ID 98111

MICROSOFT RESOURCES

"Open Source at Microsoft: World of Choice," www.microsoft.com/opensource/choice.mspx

"Interoperability Principles," www.microsoft.com/interop/principles/default.mspx

application. Then, you'll be able to bring those together. So think about that from a servicing model. If you need to patch the OS, you only have to patch a handful of images because the applications are separate. If you need to patch an application, you patch the single application and don't have to update the OS.

Forster: What are the implications for System Center?

Anderson: At MMS, we're demonstrating the combination of Microsoft Application Virtualization version 4.5 and how it integrates with the SCCM 2007 R2 release. With Softricity, you package—we call it "sequence"—the application. Now from within the Sequencer, you'll be able to automatically publish that application into SCCM, just like you can any other MSI application. SCCM will automatically replicate those applications around the world, down to your end devices, and put them in your cache. So when the end users click on the application, they won't know if the application is virtual or standard. The other thing we've done is taken the ability to stream down the bits required to get an application running and then bring the other bits. This streaming server now becomes a role inside SCCM.

Forster: How do you summarize the focus of MMS this year?

Anderson: There are all these different computing models and ways of accessing data and applications. System Center can give you one consistent way to get to your applications and data, the things you need to do your work, anytime, anywhere. And we'll manage all that intelligently. The goal is a consistent working experience, independent of device or location.

The Big Picture

Microsoft is moving from fear of open-source competition to a newfound confidence in the value proposition of a consistent and unified management infrastructure across servers, clients, and applications. (For Anderson's and Oreck-lin's perspective on the "Dynamic Desktop" aspect of the new strategy, see the Web-exclusive sidebar "System Center and Anytime, Anywhere, Any-Device Management," InstantDoc ID 98434.) This revolution has been developing for the past couple of years. But the breakthrough came when Microsoft realized that virtualization creates an opportunity to embrace competition while defining a competitive edge based on a unified management infrastructure. Combine that with emphasis on playing nicely with (some) competitors, Software + Services, a platform orientation, and reliance on established Microsoft skill sets, and the company looks to have a new determination and energy to dominate the market.

InstantDoc ID 98432

Karen Forster
(karen@windowsitpro.com) is group editorial and strategy director for *Windows IT Pro* and *SQL Server Magazine* and former director of Windows Server User Assistance at Microsoft.

Managing PST Files in Microsoft Outlook

Personal folders have evolved into a flexible and stable email-storage mechanism

by William Lefkovics

ILLUSTRATION BY IMAGEZOO/IMAGES.COM

When Microsoft Outlook doesn't depend on Exchange Server for email services, it uses personal folder (PST) files to store mailbox data. Also, where Exchange profiles are in use and Exchange Cached Mode (previously called offline mode) is configured, OST files are stored on the client. An OST file is a slave copy of a specific Exchange mailbox but otherwise maintains the same properties of a PST file. In her article "Common .pst File Questions" (InstantDoc ID 24017), Sue Mosher discussed some of the concerns and solutions associated with managing PST files for users. Most of those concerns remain valid today, but there are some new ones that warrant your attention. To make sound decisions about PST file usage in your environment, you need to understand a little bit about the evolution of PST files in general, as well as the differences between legacy PST files and the newer Unicode-based PST files. You might need to know how to manipulate their location and use according to your unique business requirements.

Personal Folder Evolution

PST files have undergone growing pains over the years. Older PST files had significant limitations that tested users' patience. Early PST files had limitations of 16,384 folders and 16,384 items per folder. Outlook 98 introduced the *Allow upgrade to larger tables* setting, which made an irreversible change to the PST file headers and increased the maximum item count per folder to 65,536.

With Outlook 2003 came the option to use a new PST format based in Unicode. Unicode, as the name suggests, provides a single code set in which unique numbers are assigned to any letter or character, regardless of language. This feature provides better support for foreign languages, including Scandinavian languages. The most notable feature of Unicode PST files compared with their ANSI predecessors is the elimination of the 2GB file size limitation, which I'll discuss later. Figure 1 shows the option to create a legacy ANSI PST file or the new Unicode PST file in Outlook 2003 on Windows XP.

Figure 1: Choosing between an ANSI PST file or a Unicode PST file

In Outlook 2003, Microsoft made many improvements to PST files, affecting both the ANSI format and the new Unicode PST format. The system now uses multiple indices and larger caches, accessing more content from a cache, including sorting choices. Unicode PST files also benefit from a larger sort buffer, which improves performance and index recoverability. In Outlook 2003 and later, PST files also reduce the amount of file-level fragmentation by increasing PST size in sensible blocks of 2MB each when the PST file is 20MB or larger. (Outlook assumes blocks of 256KB for PST files smaller than 20MB.) This *just in time and just enough* algorithm makes for more efficient use of drive space. And the system has a better handle on managing table indices to free up resources. For example, Outlook will no longer store indices for smaller tables when they can be easily recreated on demand. PST files created with Outlook 2003 or later recycle item IDs, letting Outlook go beyond the previous limitation of

128 million items over the life of the PST file. However, a PST file can still have a maximum of only 128 million items at one time. Hopefully, you haven't hit that ceiling!

In addition to offering full Unicode support, Unicode PST files have the benefit of removing the crippling 2GB limitation that legacy PST files face. Perhaps you understand the frustration of hitting the 2GB threshold without reasonable warning and without reasonable or obvious options once you've hit it. When trying to access an ANSI PST file that has reached the critical 2GB size, Outlook will return an error message such as *Errors have been detected in this file. Quit all mail-enabled applications, then use the Inbox Repair Tool.* The Inbox Repair Tool (scanpst.exe) won't resolve this error. As outlined in the Microsoft article "Oversized PST and OST crop tool" (support.microsoft.com/kb/296088), the data storage file has exceeded the 2GB limit and needs to be cropped. The article provides a link to the Microsoft Download Center, from which you can download a utility called Pst2gb (pst2gb.exe). Pst2gb essentially trims the PST file to a usable size without user input on what content gets unceremoniously removed. The new limit is 33TB—more theoretical than practical. (I don't think you'll need to call Microsoft anytime soon for the Pst33tb.exe utility.) And because Unicode PST files also shed the limit of 65,536 items per folder, there's no limit to the number of items per folder except for the 128 million item limit for the entire PST file, as I mentioned earlier.

Converting to Unicode

The new benefits of Unicode PST files are automatically realized when you create them in Outlook 2007 or Outlook 2003. However, existing PST files in ANSI format represent an administrative challenge should they require upgrading to the new Unicode PST format.

Learning Path

WINDOWS IT PRO RESOURCES

"Common .pst File Questions," InstantDoc ID 24017

"Dealing with .pst Files," InstantDoc ID 40961

"Your .pst File Questions Answered," InstantDoc ID 47196

"Managing Personal Folders," InstantDoc ID 50442

You can convert existing legacy PST files to Unicode PST files in several ways.

Although there's no in-place upgrade for actual PST files, you can use a few methods to move PST content from ANSI PST files to new Unicode PST files manually. First, when you create a new profile without pointing to a specific PST file, Outlook will create a new Unicode PST file in the default location. At this point, you can either use File, Import and select the ANSI PST file or open the ANSI PST file from within Outlook and manually move the data to the new PST file. (Whereas the Import option changes the time stamp on the data to the date and time of the import, manually moving the content between PST files retains the original time stamp.) Second, you can use Outlook's archive feature to move the data. Create a new Unicode PST file (be selecting File, New, Outlook Data File), then select File, Archive and select the new Unicode PST file as the archive destination. When the data has been archived, select the new Unicode archive PST file as the destination for new content for the profile.

You can also use a third-party application to accomplish this task. For example, Pete Maclean's Upstart (www.maclean.com/upstart/index.html) can migrate content from a set of PST files in a common location to new Unicode PST files, leaving the originals intact. This solution might make particular sense for enterprise migrations of PST data in which large numbers of PST files need conversion to Unicode PST files. An enterprise solution might save a lot of time from manually migrating large numbers of PST files.

No matter which method you use, you should save a backup copy of the original PST files before moving content. Converting to Unicode PST files isn't a reversible process, and the result is that Outlook versions prior to Outlook 2003 won't be able to access content. Trying to use Outlook 2002 or earlier to open Unicode PST files will return an error such as *This .pst file is not compatible with this version of the Personal Folders information service. Contact your Administrator.*

Are PST Files Bad?

There's a long-running mantra among Exchange administrators that PST files are "bad." What those administrators mean is that, for many reasons, a PST file isn't the best choice as a primary storage location for mailbox data. Consider the following:

- The Microsoft article "Personal folder files are unsupported over a LAN or over a WAN link" (support.microsoft.com/kb/297019/en-us) advises that PST file access isn't supported over the network.
- PST files need to be stored locally on users' workstations, which can complicate your backup and recovery efforts.
- PST files consume more drive space than the equivalent content on the Exchange server.
- PST files have no mechanism to support Single Instance Storage (SIS).
- Content stored in a PST file can't be accessed simultaneously from another source, such as another workstation or Outlook Web Access (OWA).
- PST files lack significant security. Passwords, for example, can be easily broken or stripped with the right tools.
- If a user needs a local copy of his or her mailbox, cached Exchange mode is preferred. This mode creates a slave copy of the mailbox using an .ost extension. The user can work with the local copy of the mailbox in the event of connectivity interruption between the workstation and the Exchange server. The OST file will synchronize with the Exchange mailbox when connectivity is restored.

Certainly, however, there are times where PST files are beneficial in a corporate environment. Remember that many businesses are using Outlook to access messaging solutions other than Exchange, and these solutions most likely employ Internet protocols POP3 or IMAP4. The latter protocol synchronizes with the server mailbox, but the former might be best configured to leave a copy of messages on the server to reduce all-or-nothing dependency on the local PST file. PST files do make inexpensive archives, but from a corporate perspective, archives might need to be centrally managed and searchable for compliance and discovery reasons.

Controlling the Default Folder Location

By default, Outlook keeps PST files in a folder deep inside the users' profile folder hierarchy, thus providing basic security when other users share the same workstation. In XP, you'll find the files at \Documents and Settings*username*\Local Settings\Application Data\Microsoft\Outlook, and in Win-

dows Vista, you'll find them at \Users*username*\AppData\Local\Microsoft\Outlook.

If you don't like email storage to be buried in a user profile, or if you use automation across different client OSs, one recommendation might be to standardize the PST file location to a specific place (e.g., D:\email). Doing so might simplify an automated backup of workstation email, for example. To configure an Outlook data storage location, you can use the Office 2007 Customization Tool (OCT), which can assign the default folder location for new PST or OST files for Outlook 2007. The tool's predecessor, the Custom Installation Wizard (CIW) for Outlook 2003, can also perform this function for new installations. Figure 2 shows the OCT option. You'll find the setting under Features, Modify User Settings. Navigate to Microsoft Office Outlook 2007, Miscellaneous, PST Settings. This configuration won't move PST files already in use but will assign a default location for their creation.

Figure 2: Standardizing the PST file location

Alternatively, to change the default path that Outlook uses when creating a new PST file, you can create a registry entry specific to the user. This value works for Outlook 2007, 2003, and 2002. Open your registry editor, and navigate to HKEY_CURRENT_USER\Software\Microsoft\Office*version number*\Outlook. (The version number for Outlook 2007, 2003, and 2002 and are 12.0, 11.0, and 10.0, respectively.) Add a new string value called ForcePSTPath with the full path as the value. A registry file for this entry for Outlook 2007 would resemble

```
[HKEY_CURRENT_USER\Software\Microsoft\
  Office\12.0\Outlook]
"ForcePSTPath"="C:\\email"
```

When creating a new Outlook Data File, Outlook will open this folder location for the user to save his or her new PST file. This location also applies to new OST files.

Configure a Maximum PST and OST File Size

Although Unicode PST files have a theoretical maximum size of 33TB, they still have a default size limitation of 20GB. You can amend this threshold lower or higher through the registry. Certainly, larger PST files aren't going to perform as well; even 20GB is quite excessive. An administrator or user might want to prevent uncontrolled growth of personal folders on client workstations because of backup and restore resource requirements, prevention of performance degradation, or even disk space concerns. My recommendation would be to limit the size to around 10GB–12GB. To do so, you can configure a user setting at HKEY_CURRENT_USER\Software\Microsoft\Office*version number*\Outlook\PST or a policy setting at HKEY_CURRENT_USER\Software\Policies\Microsoft\Office*version number*\Outlook\PST. Again, the version number applies to the version of Outlook (and Office).

You can create four entries for each of these registry locations. The DWORD values that control this setting for Unicode PST files are MaxLargeFileSize and WarnLargeFileSize. The latter value doesn't actually warn the user, but it does let Outlook continue with internal operations while preventing the arrival or creation of additional content. (The PST file can still marginally expand as a result of internal processes when the WarnLargeFileSize value is reached.)

ANSI PST files have a similar control but with smaller values. Because ANSI PST files still have the 2GB limitation, their registry entry is in bytes and the word *Large* is removed from the DWORD value names. If the MaxFileSize for an ANSI PST file is set beyond the 2GB limit, the value is ignored.

An example of these registry settings for an Outlook 2007 user with the default values would resemble

```
[HKEY_CURRENT_USER\Software\Microsoft\
  Office\12.0\Outlook\PST]
"MaxLargeFileSize"=dword:00005000
"WarnLargeFileSize"=dword:00004c00
"MaxFileSize"=dword:7bb04400
"WarnFileSize"=dword:74404400
```

The above DWORD values are in hexadecimal format. You can easily push out these registry entries through Group Policy, the *Microsoft Office Resource Kit* tools, logon scripts, or another remote-management application.

Disable PST/OST Creation

Sometimes in an Exchange environment, you actually want to prevent users from using PST files. Perhaps you have a server-side archive solution, or access to workstations is too insecure for email storage; in such cases, you might require that all content remain in a centralized location. Again, you can accomplish such a configuration through a registry entry that you can then push to systems through Group Policy, logon scripts, the OCT for Outlook 2007, or the CIW for new Outlook 2003 installations. You could also use your registry editor to manually perform the configuration.

Navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office*version number*\Outlook, and create DisablePST—a DWORD entry with a decimal value of 1 to toggle it on. This value removes the Outlook Data File option from the File, New menu in Outlook, as well as the File, Archive option. It also prevents the creation of OST files and archive PST files and disallows exporting to or importing from other PST files.

Flexibility and Stability

Newer Unicode PST files bring greater flexibility and stability to managing messaging data with Outlook 2007 and Outlook 2003. Different companies will have different needs regarding email storage, and PST management might be part of that administrative challenge. Remember that you have options available to you for controlling PST file location and size—and even for preventing their use.

InstantDoc ID 98540

William Lefkovics
(william@mojavemediagroup.com) is a technical writer specializing in messaging and collaboration solutions and is technical director of Mojave Media Group. He is an MCSE and a Microsoft Exchange MVP.

Use OpenLDAP's proxy service to allow

Integrate Active Directory and OpenLDAP

BY DUSTIN PURYEAR

SOLUTIONS SNAPSHOT

PROBLEM:
You can't access Active Directory (AD) Schema via OpenLDAP.

SOLUTION:
Use OpenLDAP's proxy service to connect to AD.

WHAT YOU NEED:
CentOS; OpenLDAP; AD running on Windows Server 2003 R2

DIFFICULTY:
●●●◒○

SOLUTION STEPS:

1. Start slapd.
2. Configure slapd-ldap; restart slapd and run ldapsearch.
3. Install OpenLDAP 2.3.
4. Modify pidfile and argsfile.
5. Restart slapd and run ldapsearch again.

Both Active Directory (AD) and OpenLDAP play important roles in the enterprise. In fact, within the same company you'll find the UNIX group using OpenLDAP and the LAN and Windows administrators using AD. However, most people are unable to fully access the AD schema via OpenLDAP.

OpenLDAP and AD can peacefully coexist—the key is finding the best way to allow LDAP operations to cross the boundaries between AD and OpenLDAP deployments. One way to make that happen is to use OpenLDAP's proxy service. To demonstrate this proxy service, I'll walk you through the steps to make AD's cn=Users container, which by default contains all user objects, part of an OpenLDAP directory.

Terms and Versions

Before moving on, let's define terminology. First, an LDAP server is actually what is known as a Directory Service Agent (DSA). Second, a DSA manages either part or all of a Directory Information Tree (DIT). Several DSAs may be deployed to manage an entire DIT as well as to allow for replication and high availability. The portion of the DIT that a DSA manages is known either as a partition or database. I use the term database.

To produce the examples in this article, I used CentOS 4.3, OpenLDAP 2.2.13, and AD running on Windows Server 2003 R2. Later in the article, I'll show you a limitation in the commonly deployed OpenLDAP 2.2, which you can solve by installing OpenLDAP 2.3 on CentOS 4.3.(For CentOS 4.3, I use the RPMS found at dev.centos.org/centos/4/testing/i386/RPMS/.) See the sidebar "Upgrading OpenLDAP on CentOS," for installation instructions.

Starting the OpenLDAP Server Process

The OpenLDAP server process is named slapd, which stands for "stand-alone LDAP daemon." It provides almost all of the OpenLDAP server functionality, including the ability to accept connections from LDAP clients,

Listing 1: slapd.conf

```
# Import our schema
include /etc/openldap/schema/core.schema
include /etc/openldap/schema/cosine.schema
include /etc/openldap/schema/inetorgperson.schema
include /etc/openldap/schema/nis.schema

# Support both LDAPv2 and LDAPv3
allow bind_v2

pidfile /var/run/slapd.pid
argsfile /var/run/slapd.args

loglevel 1

# Our primary back end
database bdb
suffix "dc=testcorp,dc=com"
rootdn "cn=manager,dc=testcorp,dc=com"
rootpw "password"
directory /var/lib/ldap
# Indexes for this back end
index objectClass eq,pres
index ou,cn,mail,surname,givenname eq,pres,sub
index uid eq,pres,sub
```

LDAP operations to cross boundaries

process queries and updates, and implement the ACLs that restrict access to confidential information within the directory. Notably, in OpenLDAP, replication is handled by another process entirely and is beyond the scope of this article.

Let's start off with a sample slapd configuration that brings up a basic DIT with no ACLs or any other special capabilities. On the OpenLDAP server, configuration starts with the slapd.conf file shown in Listing 1. In this configuration, slapd manages a database for the directory tree dc=testcorp,dc=com.

To start slapd, type the following:

```
# service ldap start
```

and load the initial entries into the database.

To load the entries, first enter the information from Listing 2 into a file named dir.ldif. These entries will define a very simple tree which has a suffix (aka root) of dc=testcorp, dc=com and two branches that are ou=People and ou=Groups. Now, load the entries using ldapadd:

```
# ldapadd -x -h localhost \
-D cn=manager,dc=testcorp,
dc=com -W \
-f dir.ldif
Enter LDAP Password: <value-of-rootpw>
adding new entry "dc=testcorp,dc=com"
adding new entry "ou=People,
 dc=testcorp,dc=com"
adding new entry "ou=Groups,
 dc=testcorp,dc=com"
```

The -x option specifies that ldapadd should use simple authentication instead of Simple Authentication and Security Layer (SASL). With simple authentication, the LDAP client (in this case, ldapadd) sends the credentials in plaintext. Even if you use LDAP over SSL (LDAPS) or LDAP StartTLS, you're still using simple authentication, but the tunnel being used for communication is encrypted (and far more secure).

We can test that our entries loaded properly by using ldapsearch

```
# ldapsearch -LLL -x -h localhost \
-b 'dc=testcorp,dc=com'
```

which performs a query to find all entries below the root of the tree. Figure 1, page 48, shows the results. As expected, ldapsearch returns the three entries that we originally imported via ldapadd. We are now ready to begin working with referrals.

Listing 2: Entries in dir.ldif to Define a Simple Tree

```
dn: dc=testcorp,dc=com
objectClass: top
objectClass: organization
objectClass: dcObject
dc: testcorp
o: Test Corp, Inc.

dn: ou=People,dc=testcorp,dc=com
objectClass: top
objectClass: organizationalUnit
ou: People

dn: ou=Groups,dc=testcorp,dc=com
objectClass: top
objectClass: organizationalUnit
ou: Groups
```

A Caveat to Using Referrals

You saw how easy it is to view entries that OpenLDAP manages by using a simple ldapsearch command on our client—but what about viewing entries that AD manages? For that to happen, you need to direct either the LDAP client or the LDAP server (i.e., OpenLDAP) to AD.

An obvious choice is to use referrals, which is a way for a DSA to forward—or refer—an LDAP request to another DSA. However, while referrals are both powerful and flexible (both for managers and application developers), keep in mind an important caveat: How a client handles a referral is entirely dependent on implementation. For example, OpenLDAP's ldapsearch can chase referrals when used with the -C option, but only anonymously—ldapsearch doesn't try to authenticate against the second DSA.

If you did create a referral in OpenLDAP to AD, ldapsearch (as well as other OpenLDAP binaries such as ldapadd) would return output containing the following: "In order to perform this operation a successful bind must be completed on the connection." This statement simply means that ldapsearch chased the referral to a domain controller (DC) and the operation was rejected because ldapsearch didn't try to authenticate.

To solve this problem, you could write your own LDAP clients (e.g., using Perl's Net::LDAP), use another vendor's toolset, or bypass this issue entirely and use OpenLDAP's proxy service. I'll show you the third option.

Upgrading OpenLDAP on CentOS

To upgrade from OpenLDAP 2.2 to OpenLDAP 2.3 on CentOS 4.3, you need to update yum, the automatic updater/package installer/remover for RPM systems. Then you're ready to upgrade your OpenLDAP packages. First, point yum at the CentOS development repository, or repo, by adding these lines to /etc/yum.repos.d/CentOS-Base.repo:

```
[testing]
name=CentOS-$releasever - Testing
baseurl=http://dev.centos.org/centos/$releasever/testing/$basearch/
enabled=1
gpgcheck=1
gpgkey=http://dev.centos.org/centos/RPM-GPG-KEY-CentOS-testing
```

Next, upgrade your OpenLDAP packages using this command:

```
# yum -y install openldap.i386 openldap-clients.i386 openldap-devel.i386
  openldap-servers.i386
```

Upgrading may take a while, as yum needs to update all packages related to OpenLDAP as well.

InstantDoc ID 98450

```
dn: dc=testcorp,dc=com
objectClass: top
objectClass: organization
objectClass: dcObject
dc: testcorp
o: Test Corp, Inc.

dn: ou=People,dc=testcorp,dc=com
objectClass: top
objectClass: organizationalUnit
ou: People

dn: ou=Groups,dc=testcorp,dc=com
objectClass: top
objectClass: organizationalUnit
ou: Groups
```

Figure 1: Output from ldapsearch

Using OpenLDAP as a Proxy

OpenLDAP can provide a proxy for connections to AD on behalf of a client. Simply put, OpenLDAP will work with AD for you whenever necessary. The benefit of this approach is that you don't have to rely on the behavior of the LDAP client—the server, OpenLDAP, will be configured to chase referrals for you so that you don't have to depend on client behavior, which may or may not work as you want.

With an OpenLDAP proxy, all operations are routed through slapd, even though some need to be performed within AD. For routing through slapd to work, you configure slapd-ldap, the proxy back end for the slapd daemon. You could use other back ends such as slapd-meta, which provides even more features such as naming context rewriting, but slapd-ldap is the simplest to configure during initial testing.

Let's add a slapd-ldap configuration to slapd.conf. Listing 3 shows the code. In it, you can see several changes:

database ldap. We have defined a new back end by using slapd-ldap, which will be our proxy service.

subordinate. Without this keyword, slapd searches only the database specified by the search base (e.g., if dc=testcorp,dc=com were the specified search base, then cn=users,dc=testcorp,dc=com would never be examined because it's a different slapd database).

rebind-as-user. This option tells slapd to bind to the remote DSA with the credentials supplied by the client; the credentials must be valid in AD.

Uri. This specifies the remote LDAP server, which in this case is the AD DC. Notice that we aren't using SSL here—in the real world, you would configure SSL for security.

chase-referrals. This option specifies that slapd will chase any referrals automatically.

What's interesting is that AD is using the same suffix of dc=testcorp, dc=com as OpenLDAP. Often both UNIX and AD administrators bring up a directory service with the same standard suffix (i.e., naming context), and only later do they find that they need to provide for better integration.

Now, restart slapd and run ldapsearch again:

```
# ldapsearch -x -h localhost
  -LLL \
-b dc=testcorp,dc=com \
-D cn=dpuryear,cn=users,
 dc=testcorp,dc=com -W \
'(cn=dpuryear)' cn
Enter LDAP Password:
```

In this example, the ldapsearch command searches against cn=users,dc=testcorp, dc=com, which slapd should map to CN=Users in AD. And slapd does map it, as you can see by the output below:

```
dn: CN=dpuryear,CN=Users,DC=t
  estcorp,DC=com
cn: dpuryear
```

which shows the entry CN=dpuryear, CN=Users, DC=testcorp, DC=com—the account for dpuryear in the AD CN=Users container. slapd now knows that any operation against cn=users,dc=testcorp,dc=com (which is our superior in the directory tree), actually requires these steps:

1. Open an LDAP connection to ldap://dc1.testcorp.com/.
2. Bind with the credentials supplied by the client.
3. Perform the operation.
4. Return the results to the client.

Listing 3: Adding a slapd-ldap Configuration to slapd.conf

```
# Import our schema
include /etc/openldap/schema/core.schema
include /etc/openldap/schema/cosine.schema
include /etc/openldap/schema/inetorgperson.schema
include /etc/openldap/schema/nis.schema

# Support both LDAPv2 and LDAPv3
allow bind_v2

pidfile /var/run/slapd.pid
argsfile /var/run/slapd.args

loglevel 1

# Our slapd-ldap back end to connect to AD

database ldap
suffix "cn=users,dc=testcorp,dc=com"
subordinate
rebind-as-user
uri "ldap://dc1.testcorp.com/"
chase-referrals yes

# Our primary back end

database bdb
suffix "dc=testcorp,dc=com"
rootdn "cn=manager,dc=testcorp,dc=com"
rootpw "password"
directory /var/lib/ldap
# Indexes for this back end
index objectClass eq,pres
index ou,cn,mail,surname,givenname eq,pres,sub
index uid eq,pres,sub
```

Listing 4: Modifying pidfile and argsfile

```
# Import our schema
include /etc/openldap/schema/core.schema
include /etc/openldap/schema/cosine.schema
include /etc/openldap/schema/inetorgperson.schema
include /etc/openldap/schema/nis.schema

# Support both LDAPv2 and LDAPv3
allow bind_v2

pidfile /var/run/openldap/slapd.pid
argsfile /var/run/openldap/slapd.args

loglevel 1

# Our slapd-ldap back end to connect to AD

database ldap
suffix "cn=users,dc=testcorp,dc=com"
subordinate
rebind-as-user
uri "ldap://dc1.testcorp.com"
chase-referrals yes

# Our primary back end
#

database bdb
suffix "dc=testcorp,dc=com"
rootdn "cn=manager,dc=testcorp,dc=com"
rootpw "password"
directory /var/lib/ldap
# Indexes for this back end
index objectClass eq,pres
index ou,cn,mail,surname,givenname eq,pres,sub
index uid eq,pres,sub
```

But there's a problem: We don't have access to all of the data in AD. To see what I mean, try to return the attribute sAMAccountName, which is specific to the AD schema. When I type

```
# ldapsearch -x -h localhost -LLL -b
  dc=testcorp,dc=com \
-D cn=dpuryear,cn=users,dc=testcorp,
  dc=com -W \
'(cn=dpuryear)' cn sAMAccountName
Enter LDAP Password:
```

INSIDE

A Security Reader

Brought to you by HELLOSECUREWORLD.COM

www.hellosecureworld.com

Network Access Protection in Windows Server 2008

VERIFY COMPUTERS' SECURITY BEFORE ALLOWING NETWORK ACCESS

By Damir Dizdarevic

[**Editor's Note:** This article is based on Windows Server 2008, post-Beta 3 June Community Technology Preview (CTP) build, which was the version available at article submission time. Note that some of the options and screens will change in the final release.]

In considering LAN security, we mostly think about preventing an attacker from accessing network resources. The reason for this focus is simple: Most attacks are initiated from the Internet and are directed at breaking into private networks.

However, an equally large security issue that administrators must address is preventing regular (i.e., authenticated and authorized) network users from using computers with weak security configurations to access network resources. For example, a traveling employee might have a laptop that only occasionally uses the VPN to connect to the corporate LAN—but this laptop still needs all the current security fixes, antispyware, and antivirus definitions installed. Otherwise, such a computer is a likely source to spread viruses or worms on the network. If a computer doesn't have a firewall enabled and becomes infected with Trojan-like software, the computer can provide unauthorized persons with easy access to local network resources. Employees' home computers that use the VPN to access the corporate LAN and that aren't managed properly provide a similar risk. Finally, letting visitors connect their computers to your local network, even just to provide them with Internet access, can put other hosts on the network at risk for infection with viruses or other kinds of malicious code.

The question is: How do you check a computer's security before you allow it to access network resources? In addition, how do you determine whether to grant full or limited access? Network administrators need a mechanism to ensure that any computer connecting to the corporate network meets the organization's health policy requirements and has all the necessary software, patches, and hotfixes installed.

Network Access Protection

Windows Server 2003 SP1 includes a technology called Network Access Quarantine (NAQ) that helps administrators limit or deny connections to computers that don't comply with a company's security policies. However, NAQ has many disadvantages. First, it's limited to VPN-based connections only, which means you can't protect your network from unsecured wireless users, or even from users who have a physical connection to the network (e.g., via employees' personal laptops). Second, NAQ is based on manually created scripts (implemented via Connection Manager) that must be run on the client side before VPN access is granted. These scripts check such things as the firewall state, antivirus state, presence of a password-protected screen saver, and status of Internet Connection Sharing. Besides the fact that writing the scripts can be difficult and time consuming, the various types of protection software on the client side can also cause problems. For example, if VPN clients have various antivirus programs, you must write a specific script for each program and use a different Connection Manager package for each. In the end, this solution is static. After the client passes all the checks and the main script reports the state of the client's health to the server, the user can safely disable the firewall, antivirus software, and all other security features. These actions won't be detected, and the level of access to resources will remain unchanged.

Windows Server 2008 solves most of NAQ's disadvantages with Network Access Protection (NAP) technology. Using NAP, an administrator can enforce specific compliance policies that must be met before a client computer can access network resources. If a client computer doesn't meet the defined health requirements, it's either placed in quarantine (with access limited to specific hosts) or simply not allowed access.

In addition, NAP can automatically remediate unhealthy clients, updating systems when

Using NAP to Verify a Computer's Security Before Allowing Network Access

1. Prepare the environment: Install DHCP on a Windows Server 2008 DC with AD enabled, and add the Network Policy Server and DHCP Server roles.
2. Configure health policies (i.e., the System Health Validator and Health Policy options).
3. Create network policies for NAP: Configure policies for incompliant and compliant clients, as well as a policy for NAP non-capable clients.
4. Configure DHCP for NAP: Configure the DHCP server to distribute a different group of scope options to compliant and incompliant NAP clients.
5. Enforce NAP on the client side: Use the NAP Client console, Group Policy, or Netsh to configure the client to work with NAP.

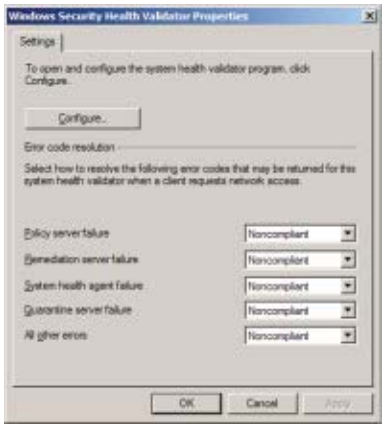

Figure 1: *Configuring Windows Security Health Validator properties*

possible to make them comply with corporate policy. The administrator configures NAP's method of enforcement, depending on the type of client connection. NAP enforces health requirements for the following types of connections:

- IPsec-protected communications
- IEEE 802.1x-authenticated connections
- VPN connections
- DHCP-managed connections

In this article, I focus on NAP implementation for DHCP-managed connections. Using NAP with DHCP lets you protect your network from all potentially unsecured clients that are managed via DHCP (i.e., clients that receive IP addresses from DHCP), including resident desktop computers that are NAP capable.

NAP-capable OSs include Windows Vista (by default) and Windows XP SP2 with NAP client software (currently in Beta 3). XP SP3 will include the NAP client by default. No older OSs are supported, because NAP relies on information from Windows Security Center (WSC), which exists only in Vista and XP SP2.

A benefit of NAP is that it's not limited to Microsoft technologies. Any system that can provide the NAP server with its health state can also use NAP. Microsoft is working with many hardware and software vendors and other partner companies to help them create NAP-compatible devices and software. To use NAP for DHCP-managed connections, you must prepare the environment, configure health policies, create network policies for NAP, configure DHCP for NAP, and enforce NAP on the client side.

Prepare the Environment

First, you must have an existing Active Directory (AD) infrastructure available, with one or more Windows 2003-based (or Server 2008-based) DCs. DHCP must be installed on a Server 2008 machine, because previous versions of the DHCP service (such as the version on Windows 2003) aren't aware of NAP. You need at least one static IP address for this host.

Install Server 2008 as a member server in your domain. After installation, you must add the Server 2008 roles called Network Policy Server and DHCP Server. You can easily accomplish this task through the Server Manager console, which is available on the Welcome page or under Administrative Tools. Open Server Manager, go to Roles Summary, and click Add Roles. Server 2008's Network Policy Server role replaces Windows 2003's Internet Authentication Service (IAS). Thus, Network Policy Server (NPS) lets you create various types of policies, not just those related to NAP.

Configure Health Policies

To configure your health policies, go to Administrative Tools and click the Network Policy Server role you added. In the NPS console that opens, you must configure the System Health Validator and Health Policy options to create an appropriate network policy. The System Health Validator component defines your security requirements for clients that are accessing the network, whereas Health Policy defines different configurations for NAP-capable clients.

Double-click the Network Access Protection node on the left side of the console, and click System Health Validator. The Windows Security Health Validator item will appear on the right side of the console. Double-click this item to open the configuration window that Figure 1 shows. In this window, click Configure to see options for security requirements. As Figure 2 shows, you can simply select the appropriate check boxes to indicate what you require from clients. In Vista, you can require the firewall to be enabled, antivirus and antispyware applications to be present and current, the automatic update feature to be enabled, and current hotfixes to be installed. Similar requirements are available in XP SP2, other than the antispyware option, which isn't part of XP. For testing purposes, let's select only the firewall check box for both Vista and XP. Click OK twice to finish configuring the System Health Validator option.

To configure the Health Policy option, double-click the Policies node in the NPS console, right-click Health Policies, and select New. In the window that Figure 3 shows, enter the policy name and select what the System Health Validator (SHV) component will check.

First, let's create a policy for compliant clients. Enter compliant for the policy name, and select Client passes all SHV checks from the drop-down menu. Selecting this option means that, for a client to be considered healthy, it must pass all the requirements you configured in SHV (which in the example was only the firewall requirement). Next, select the Windows Security Health Validator check box and click OK. Your first policy is now configured.

Next, let's create a policy for non-compliant computers. Follow the same steps to create a new health policy, perhaps called noncompliant. In the drop-down menu, select Client fails one or more SHV checks, which means that if a client fails to correctly report one or more required components from SHV, it will be considered unhealthy. Finally, select the Windows Security Health Validator check box and click OK.

Create Network Policies for NAP

After you configure the SHV and Health Policy options, you can configure network policies. In the NPS console's Policies node, click Network Policies and disable the default policies. By default, the two default policies are Connections to Microsoft Routing and Remote Access Server and Connections to other access

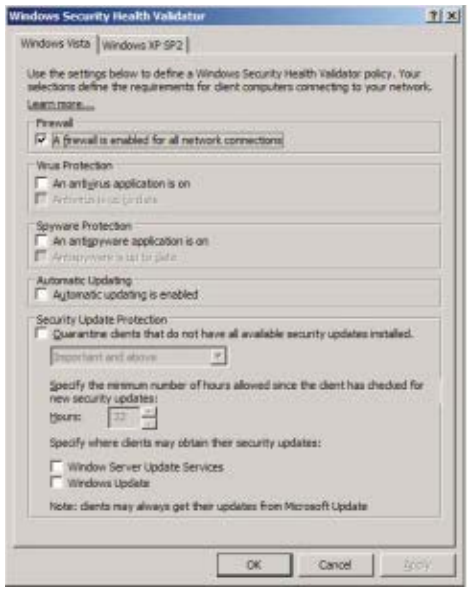

Figure 2: *Setting security requirements for Windows Vista clients*

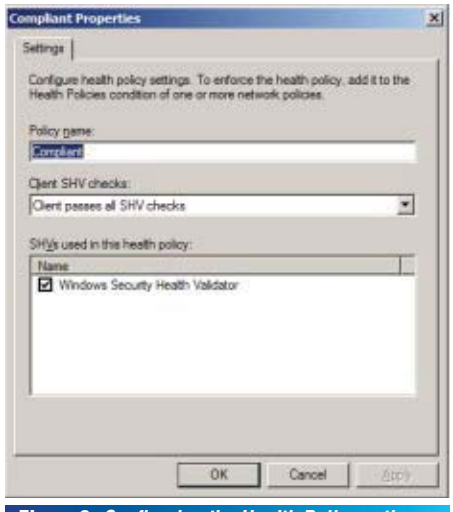

Figure 3: *Configuring the Health Policy option*

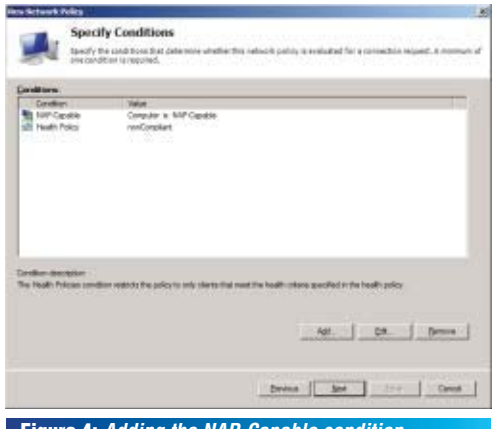

Figure 4: *Adding the NAP-Capable condition*

servers. Right-click each policy and select Disable from the drop-down menu. Then, right-click the Network Policies folder and select New. A wizard for creating a new policy will start. Enter a policy name (e.g., "noncompliant-restricted" for a policy for unhealthy clients). Then, select from the drop-down list the type of network access server that will apply the policy to clients. The default is Unspecified; for our purposes, select DHCP Server.

Click Next to proceed to the Conditions page, and click Add to select conditions for the policy. From the list of available conditions that displays, select Health Policies from the Network Access Protection group. In the window that opens, select the "noncompliant" health policy that you created earlier.

Follow the same procedure to add the NAP-Capable condition, which Figure 4 shows, to the policy. This condition limits application of the policy only to computers that are NAP capable.

Click Next to launch the Specify Access Permission window. In this window, you must specify what to do with clients that meet the policy. Although denying access to unhealthy clients might seem logical, you don't want to completely deny access to those clients. Instead, you should provide them with limited access only to hosts that can help them improve their security state (i.e., remediation servers). Select Access Granted and click Next.

In the Configure Authentication Methods window, select the Perform machine health check only option, and clear the other check boxes, as Figure 5 shows. Because you're configuring a policy for checking clients' security health state via DHCP and because DHCP clients don't authenticate to the DHCP server, you don't need to configure authentication methods. Just click Next in the Configure Constraints window—none of the options apply to our example.

In the Configure Settings window, select NAP Enforcement in the Network Access Protection section, as Figure 6 shows. For this policy, you should select the Allow limited access NAP enforcement method. This setting will put clients in quarantine and give them access only to remediation servers. You can also configure those servers from this window: Simply click Configure to create a Remediation Server Group, and enter IP addresses for the hosts. Also select the Enable auto-remediation of client computers check box. Enabling both these settings causes the NAP Enforcement client component to automatically attempt to update the computer security state (e.g., if you turn the firewall off, it will be turned on automatically).

After you create a policy for incompliant clients, you must create a policy for compliant clients. Follow the same steps to create another new network policy, this time naming the policy "compliant full." On the Conditions page, select the "compliant" health policy. Then, select the Allow full network access check box on the NAP Enforcement Settings tab. All other settings are the same as for the incompliant client policy.

Finally, you can configure a policy for NAP non-capable clients, to provide them with or deny them network access. This policy should grant or deny access to clients that aren't NAP capable, by implementing only the NAP-Capable condition, with the Only computers that are not NAP-capable option selected. (Note that this policy isn't necessary in a test environment.)

Figure 7 shows the NPS console after you've created the necessary policies. Next, you need to configure DHCP.

Configure DHCP for NAP

You need to configure DHCP, so that DHCP can use NPS and the policies you created. First, you must create a scope on the DHCP server. Our intention is to configure the DHCP

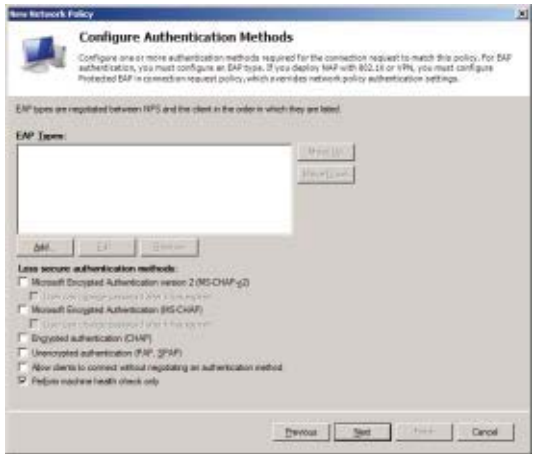

Figure 5: *Configuring a policy for checking clients' security health state via DHCP*

Figure 6: *Configuring NAP enforcement*

server to distribute a different group of scope options to compliant and incompliant NAP clients. After you create a scope, right-click it in the DHCP console. Select Properties, and go to the Network Access Protection tab. Then, select the Enable for this scope check box, as Figure 8 shows, and use the default NAP profile.

Another thing you can configure from the Network Access Protection tab in IPv4's properties is DHCP behavior, in case DHCP can't contact the network policy server. The default setting is to give clients full access, but you can also select the Restricted Access or Drop Client Packet options. In addition, you can enable and disable NAP on the server level.

Finally, you must configure additional options for NAP-capable clients. Right-click Scope Options, and select Configure Options. In the Configure Scope Options dialog box, select the Advanced tab. Select Default Network Access Protection Class as a User class, and define specific DHCP options for this class of clients (e.g., different DNS domain name).

Enforce NAP on the Client Side

The last step is to configure the client to work with NAP. In fact, you must enforce the use of NAP on clients. You can accomplish this task through the NAP Client console, Group Policy, or Netsh (which has the new context for NAP configuration). Because you can't configure domain or OU Group Policy Objects (GPOs) to include NAP settings from Windows 2003, using Group Policy requires you to edit GPOs from Vista or Server 2008's Group Policy Management Console (GPMC). Use the Administrative Tools' Services console to start the Network Access Protection Agent service, changing this service's startup type to Automatic (which you can also use Group Policy to accomplish).

On Vista, start the Microsoft Management Console (MMC) and add the NAP Client Configuration snap-in. Alternatively, select Run from the Start menu, and enter

```
napclcfg.msc
```

Select the Enforcement Clients node in the left task pane, double-click DHCP Quarantine Enforcement client on the right side, select Enable this enforcement client, and click OK. Figure 9 shows the results. From now on, the client should be able to use NAP.

To use Netsh to configure NAP on a client, go to the command line and enter

```
Netsh nap client set
enforcement ID = 79617
```

If you want to use XP SP2, you must install the NAP client software for XP Beta 3, which makes the OS NAP capable.

Run a NAP Test

To test NAP on a client, configure a Vista client and join it to your domain. Obtain an IP address from the DHCP server, with the firewall in the default active state. Ensure that you have a regular IP address, from the scope that you created in earlier steps, with regular scope options. To verify that you have

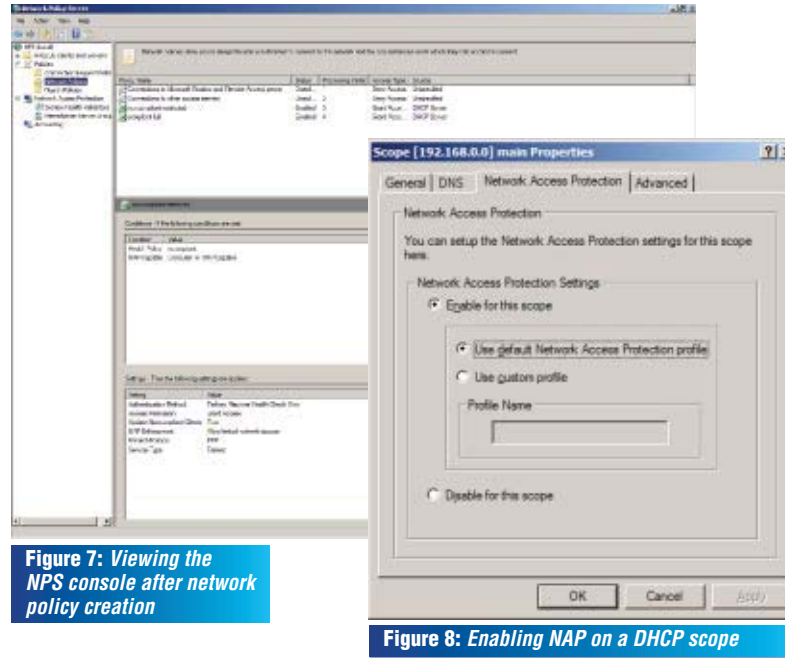

Figure 7: *Viewing the NPS console after network policy creation*

Figure 8: *Enabling NAP on a DHCP scope*

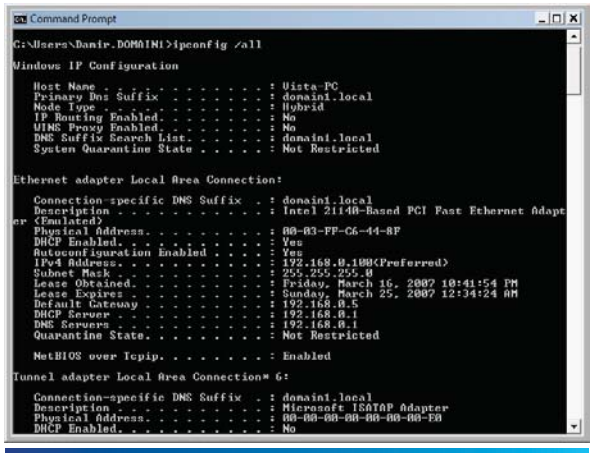

Figure 9: *Results of testing NAP on a compliant machine*

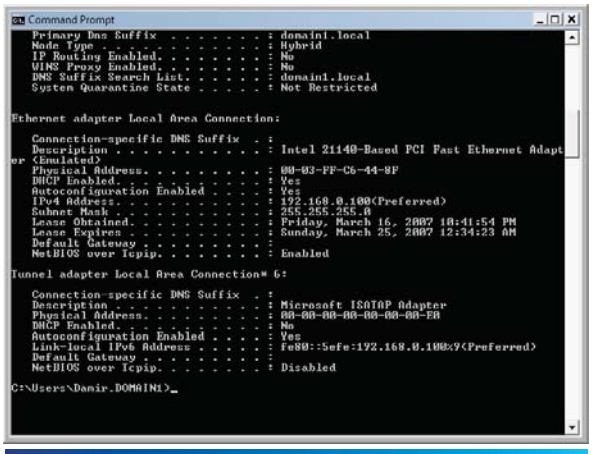

Figure 10: *Results of testing NAP on an incompliant machine*

all the necessary DHCP information (e.g., DNS servers, gateway, WINS servers), go to the command line and enter

```
ipconfig /all
```

Figure 10 shows the output.

Next, manually disable the Vista firewall. In a few seconds, the DHCP enforcement client will perform autoremediation to correct the client's system state, thus reenabling the firewall. To demonstrate a quarantined client, go to Server 2008's NAP console and configure Windows Security Health Validator to require an antivirus application to be installed and updated. If you don't have a NAP-configured antivirus solution on the Vista client, run ipconfig /release followed by ipconfig /renew to quarantine your client and receive a taskbar quarantine notification message. Run ipconfig /all again, and note that your computer is configured with the options you specified in DHCP's Network Access Protection class. As Figure 11 shows, all you have is an IP address and subnet mask—no Internet access, and no access to other hosts on the network.

An Effective Solution

Maintaining computers' health is one of the most time-consuming challenges that any network administrator faces. This complex task is made even more difficult if you must maintain system health for users who connect from home systems, partner computers and laptops that aren't under control of administrators, or computers that aren't managed through a corporate patching system (e.g., Windows Server Update Services—WSUS, Microsoft Systems Management Server—SMS). NAP is an effective solution for controlling network computers' security health.

InstantDoc ID 95617

Damir Dizdarevic *(ddamir@logosoft.ba) is the manager of the Learning Center at Logosoft in Sarajevo, Bosnia. An MCSE, MCTS, MCITP, and MCT, he specializes in Windows Server security and has published more than 350 articles in IT magazines.*

Windows without Windows

SOME BASIC COMMANDS WILL GET YOU STARTED WITH THE GUI-LESS SERVER CORE IN WINDOWS SERVER 2008

By Russell Smith

Windows Server 2008's Server Core edition is a stripped-down version of the OS—a kind of Windows lite that you control from the command line rather than from a GUI. What are the benefits of such a configuration? Its footprint is about 3MB, considerably less than a full installation of Windows Server. Of course, 3MB is just to host the OS and any server roles—it doesn't include additional data, such as Active Directory (AD) databases, that you might need for a particular server role. Server Core installs only the necessary components for any of its supported server roles. This reduces the attack surface of the OS, improves its security, and makes it easier to maintain and manage (albeit with a reduced armory of tools). New technologies in Server 2008, notably BitLocker and the read-only domain controller (RODC) functionality, can be used in combination with Server Core to provide even better security.

Are the benefits of reduced resource utilization and improved security offset by a server that some might consider hard to set up and administer? A look at the installation process and some basic configuration commands will help you get Server Core running and connected to your network so that you can begin to answer that question for yourself.

Installation and Setup

Installing Server Core is essentially the same as installing the full version of Server 2008; you simply need to select the Server Core entry instead of the Server option in the installation program (as Figure 1 shows). Not only is Server Core installation extremely simple but, as you might expect, much faster than installing the full edition of the server.

After installation has finished, you're asked to press the usual key sequence of Ctrl+Alt+Del to open the log-on dialog box. It might be a little disconcerting to then be presented with the option of logging on as Other User. Only one user is enabled by default in Server Core, and that's the administrator. Initially, no password is defined for the administrator account; you must set it the first time you log on. To do so:

1. Click Other User.
2. In the dialog box shown in Figure 2, enter administrator as the username in the upper box, and leave the lower (password) box empty. Click the arrow to the right of the boxes.
3. Enter a password.

To log off, simply type logoff at the command prompt.

Give Server Core an IP Address and Host Name

You can assign a static IP address and DNS server to a network adapter by using the netsh command, the same way you would with the full version of Server 2008. To assign an IP address, use a command like

```
netsh interface ipv4 add
address
  "Local Area Connection"
  192.168.1.100 255.255.255.0
  192.168.1.11
```

where 192.168.1.100 is the IP address, 255.255.255.0 is the subnet mask, and 192.168.1.11 is the gateway address. Of course, you should enter the full command without line breaks on the command line.

To assign a DNS server, type

```
netsh interface ipv4 set
dnsserver
  "Local Area Connection"
  static 192.168.1.101
```

where 192.168.1.101 is the DNS server's IP address.

Rename and Activate the Server

If you want to rename the server, you first need to determine the name that was automatically assigned during the installation process. Type hostname at the command prompt to return the server's name, then issue the following two commands to change the name and reboot the server:

```
netdom renamecomputer
  %computername%
  /newname:servercore1

shutdown /r
```

Once Server Core is connected to the Internet, you can activate the server by running the following command:

```
cscript c:\windows\system32\
  slmgr.vbs  ato
```

Enable Remote Desktop

Probably the two most crucial remote tools that you'll want to use with Server Core for administration initially are Remote Desktop and the Microsoft Management Console (MMC) Windows Firewall with Advanced Security snap-in. First, I'll show you how to enable and use Remote Desktop, then I'll address accessing Server Core remotely with MMC and the Windows Firewall with Advanced Security snap-in.

Although it's possible to make a Telnet connection to Server Core, Remote Desktop is the preferred method because it provides encryption, network level authentication, and other conveniences such as cut and paste. But don't get too excited—Remote Desktop won't give you a full-fledged Windows Desktop from which you can administer the server. You'll just see a command prompt as you would from the console.

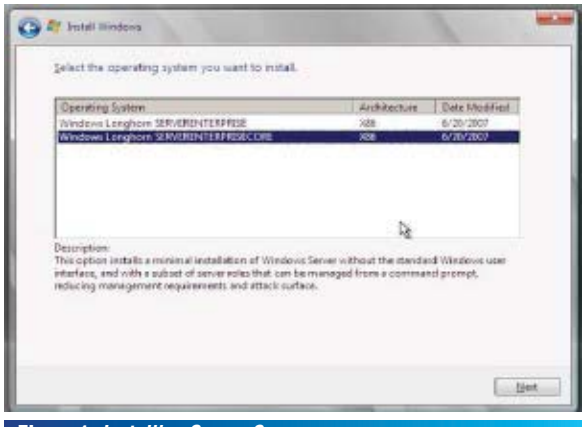

Figure 1: *Installing Server Core*

Figure 2: *Logging on for the first time*

LEARNING PATH

***WINDOWS IT PRO* RESOURCES**

For more information about Server Core:

"Sampling Server Core," InstantDoc ID 96438

"What You Need to Know About Windows Server 2008 Beta 3," InstantDoc ID 96068

Top 10: "Questions About Server Core," InstantDoc ID 96540

To learn more about RODC:

"Fortify Remote-Server Security" Security Pro VIP, InstantDoc 97962

For more information about configuring rules in Windows Firewall:

"Windows Firewall Shows New Maturity in Vista," Security Pro VIP, InstantDoc ID 95099

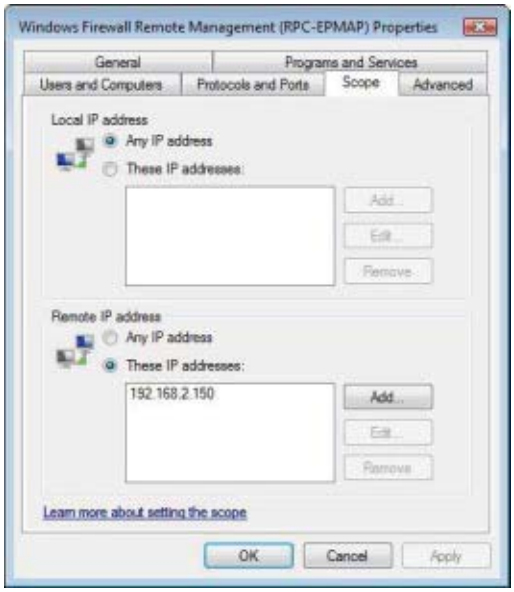

Figure 3: ***Changing the scope of an inbound Windows Firewall rule***

Because there's no command-line tool or MMC snap-in from which you can enable Remote Desktop on Server Core, you'll need to run the scregedit.wsf script that's provided as part of Server Core. Scregedit contains various functions that are the only means of performing some tasks such as setting the size of the page file, enabling Terminal Services, and product activation. To run scregedit on Server Core, use the command

```
cscript c:\windows\system32\
  scregedit.wsf /AR 0
```

If you want to access Server Core by using Remote Desktop from a Windows OS other than Vista, replace the /AR 0 switch with /CS 0. To see the full list of scregedit's possibilities, type the command

```
cscript c:\windows\system32\
scregedit.wsf /cli
```

Authenticate to Server Core with MMC

During the initial configuration, or if Server Core will be a standalone server, you might need to authenticate to it from a remote machine by using pass-through authentication. Some, but not all, MMC snap-ins let you specify a username and password when you're connecting to a remote computer.

The easiest way to get access remotely with MMC is to create a local user on Server Core that has the same username and password as the remote account that you're using to run MMC. This way, authentication will happen transparently. The new user also needs to be an administrator on Server Core to gain unrestricted access. You can create a user and add the username to the administrators group by entering the following commands:

```
net user /add
<username>
  <password>

net localgroup
administrators
  /add <username>
```

If you join Server Core to a domain, you should delete this account and use a domain-based user for authentication. Whether Server Core is a member of a domain or a standalone server, you should consider configuring Windows Firewall with Advanced Security to restrict which machines can connect remotely to Server Core.

Configure Windows Firewall

To enable the Windows Firewall with Advanced Security snap-in on any machine used for administration to access a given Server Core box, log on to Server Core as an administrator and type the command

```
netsh advfirewall set
publicprofile
  settings remotemanagement
  enable
```

To access other remote administration tools, such as the MMC Event Viewer snap-in, run the following command on Server Core to permit access through Windows Firewall:

```
netsh firewall set service
  remoteadmin enable
```

After you've made these basic changes to Windows Firewall on Server Core, you can use the Windows Firewall with Advanced Security snap-in from a remote computer for all further configuration of Server Core's firewall. You could additionally modify the firewall rules to allow access to Server Core from specific administration workstations only, if desired. To do so, you change the scope of the predefined inbound rules for Windows Firewall Remote Management, Remote Desktop, and Remote Administration by setting a list of remote IP addresses that are permitted to access Server Core. Figure 3 shows setting the scope of a Windows Firewall Remote Management rule.

Firewall rules are associated with one of three network profiles: Domain, Private, or Public. (Server Core uses the Public network profile out of the box.) To determine which profile is currently active, click the Windows Firewall with Advanced Security node directly below Console Root in the MMC window. You'll see an overview of the firewall's settings in the central pane, including information about the active profile. If you change the scope for a rule that's associated with a profile that's not currently active, the changes won't be effective.

For more informa-tion about configuring Windows Firewall with Advanced Security, see the Security Pro VIP article "Windows Firewall Shows New Maturity in Vista," April 5, 2007 (InstantDoc ID 95099). The configuration process is similar in Server 2008 and Vista.

Access the File System

The easiest way to get access to Server Core's file system is to use Windows Explorer on an administration workstation and map drives to the root administrative shares that are enabled by default on Server Core (e.g., c$ and d$). You can connect to these shares only with an account that has administrator privileges on Server Core, and you must enable remoteadmin by using netsh, as shown earlier. The File Server role is installed by default to provide access to these administrative shares, but you can also install features such as File Replication Service (FRS).

To map a network drive to an administrative share on Server

Core from a remote machine, use a command similar to the following:

```
net use z: \\192.168.1.100\c$
```

Join Server Core to an AD Domain

You can use the netdom command to join Server Core to an existing AD domain, as follows:

```
netdom add <machine name>
  /domain:<domain name>
  /userd:<user name>
  /passwordd:<password>
```

Install Server Roles and Optional Features

Server Core supports the server roles Active Directory Domain Services, Active Directory Lightweight Directory Services (AD LDS), DHCP Server, DNS Server, File Services, Print Server, Streaming Media Services, and Web Server (IIS), among others. For a full list of server roles and other supported features, go to http://www.microsoft.com/windowsserver2008/servercore.mspx.

With the exception of the Active Directory Domain Services role, you install server roles and features by using the ocsetup command. To list the server roles and features currently installed, run the oclist command. The syntax for ocsetup is the same for both roles and features. The ocsetup command-line tool is case sensitive, but you can get the correct capitalization for a server role or feature from the output of the oclist command (which Figure 4 shows). The following command installs Windows Backup:

```
start /w ocsetup
  WindowsServerBackup
```

Using the /w switch with the start command gives the user an indication of when ocsetup has finished installing the new role or feature by preventing further input at the command prompt until installation is complete. It also stops the user from running another command while ocsetup in running.

To promote Server Core to a DC, you need to generate an unattended .txt file on a full version of Server 2008 and then run dcpromo as shown below on Server Core:

```
dcpromo
  /unattend:<unattendfile.
txt>
```

Other Ways to Administer Server Core

As if these weren't enough ways to administer Server Core remotely, you can make use of Windows Remote Shell (WinRS) in Vista. The WinRS client passes commands to a WinRS listener on Server Core, which in turn passes the commands to a prompt, captures the output, and passes it back to the WinRS client. To configure WinRS on Server Core, run the following command:

```
winrm quickconfig
```

This command will prompt you to perform a couple WinRS configuration steps.

Below is an example of a command being run against Server Core remotely by using WinRS. You should note that this command line is for a machine that's a DC or domain member:

```
winrs -r:http://<servername>
  ipconfig
```

The one big disadvantage of WinRS is that it can't run commands interactively.

You can also use administration tools such as the Windows Management Instrumentation command line (WMIC) and PowerShell by means of WMI calls to manage Server Core. Unfortunately, Server Core doesn't support PowerShell directly at the time of writing (as of Server Core RC0) because PowerShell relies on the .NET Framework. Hopefully, both will be supported in a future release.

Activate Automatic Updating

You can activate automatic updating on Server Core by using scregedit to modify the registry and then restarting the Windows Update service, as follows:

```
cscript c:\windows\system32\
  scregedit.wsf /au 4
net stop wuauserv
net start wuauserv
```

As of Server Core RC0, scregedit with the /au 4 switch sets the time for checking updates to the default of 3 A.M. In Server 2008, /au 4 also reboots the server automatically if the updates require it. You can disable automatic updating by using the /au 1 switch and then restarting the Windows Update service. To check the value set for /au, use the /au and /v switches in sequence.

To force an immediate check for updates, you can use the wuauclt command as follows:

```
wuauclt /detectnow
```

Run Antivirus and Other Applications

Windows Installer is supported on Server Core, so you can use the msiexec command to install antivirus and other third-party applications if required. (And who wants to run a server without proper antivirus and backup software these days?) Before you deploy any such solution, though, you should check that it's officially supported on Server Core by contacting the vendor.

You can run at least two Windows-based applications from the console: notepad and regedit. These are useful tools, but I found it a little odd to be able to run regedit but not dcpromo, with its simple GUI.

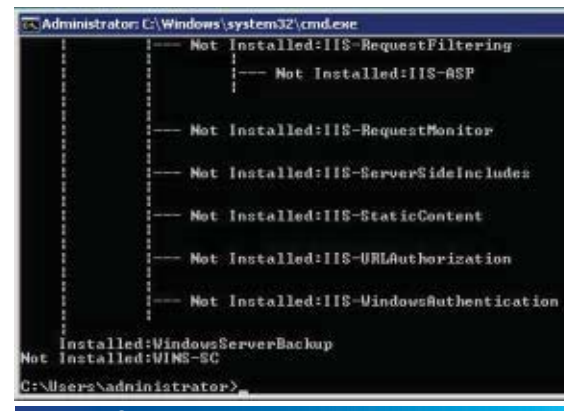

Figure 4: *Oclist command output*

Potential Not Yet Fully Realized

One of the biggest potential uses for Server Core—as a Web server—is unlikely to be realized with the current incarnation of the product due to its lack of support for the .NET Framework. Other uses, although limited, could be practical in situations that don't require frequent changes to server configuration, such as an RODC. Given the trend towards virtualization, Server Core and the hypervisor feature in Server 2008 (a software virtualization layer that sits between the hardware and the OS) together could prove to be one of the "must have" features of the next generation Windows Server.

The lack of a GUI in Server Core needn't be a disadvantage. Once the basic configuration has been completed, most other settings can be either pushed out automatically by using Group Policy if the server is part of a domain or by using MMC snap-ins on a remote machine. PowerShell is noticeable by its absence, considering it's Microsoft's latest solution for managing Windows from the command line. It's slated for inclusion in future versions of the product.

Despite some of the shortcomings of Server Core, the ability to run Windows with a significantly reduced footprint has the potential to give substantial improvements in security, capacity for virtualization, and performance.

Russell Smith *(rms45@rsitc.com) is an independent IT consultant. He has been working in IT for seven years, specializing in systems management.*

Windows Vista's Wireless Security

LET YOUR USERS GO WIRELESS WITHOUT WORRIES

By Damir Dizdarevic

Almost every time I advise someone to use a wireless rather than wired networking solution for their small office/home office (SOHO) or their home, I get a quizzical look and the inevitable question "Is that secure?" Admittedly, security is a big concern on wireless networks because wireless networks are more open to anonymous access than physical networks are. However, my typical response is that although wireless can be nonsecure, it doesn't have to be—it all depends on how much you care about security. The reality is that some people simply don't care about their computer security, perhaps because of lack of knowledge or because they think they have nothing to lose even if someone does break into their network. But if you're reading Windows IT Pro, you undoubtedly do care about security.

Windows Vista is a very wireless-friendly, as well as a very secure, OS. In this article, I explain how to use Vista's wireless networking features to enhance wireless security from the client side. These features let users configure more secure wireless networks and achieve better wireless functionality than in previous OSs.

Wireless Administration

In previous versions of Windows, hardware vendors typically provided their own tools for managing wireless networks. This method was challenging for both users and support technicians because users needed to learn how to use different vendor-specific wireless software depending on the type of computer or network adapter they had, and support personnel had to manage these various clients with different tools—mostly in a decentralized manner. Vista includes wireless client software by default. This software is hardware-vendor independent, and the interface for administering wireless networks is the same for both users and administrators. This single point of administration offers a new level of consistency for wireless clients and makes managing wireless security easier than ever before.

For additional functionality, hardware vendors and developers can use Microsoft's Extensible Authentication Protocol (EAP) architecture, called EAPHost. EAPHost is basically a framework for creating authentication mechanisms that Vista doesn't support natively. Hardware vendors or developers can use EAPHost to create a plug-in for an existing Vista wireless client, in order to provide additional authentication or encryption functionality, instead of writing a complete software package. This additional authentication functionality is available to users through the Vista wireless client (rather than in a separate application as with previous versions of Windows).

Connecting to Wireless Networks

One of Vista's most significant improvements to wireless security is that the wireless client discloses much less information about configured wireless networks. In previous versions of Windows, such as Windows XP, the client periodically broadcasts the Service Set Identifier (SSID) names of all the configured wireless networks. Malicious users can take advantage of this behavior by catching these broadcasts, then tricking a client into connecting to a false Access Point (AP), using an SSID name that matches the SSID name of a real wireless network that's configured on the client, in order to obtain private information such as a username and password for connecting to a real AP.

In Vista, a wireless client doesn't broadcast all configured SSID

names. Instead, the client broadcasts only those SSIDs that are explicitly configured as hidden and preferred networks, and only if necessary (e.g., when a user initiates a connection to configured wireless networks). If a user doesn't have any hidden networks configured, no broadcasts will occur from the client side, which greatly enhances security. (Note that using hidden SSID networks isn't a recommended practice because doing so provides only an illusion of security. Even if your AP doesn't broadcast SSID names, your clients do. Because you have many more clients than APs, and because clients are mobile whereas APs are static, a malicious user will more likely discover a hidden SSID name by sniffing client broadcast traffic rather than obtaining the name from an AP.)

Vista helps users connect to hidden networks by displaying "unnamed networks" in the Connect to a network wizard, which Figure 1 shows. To access this wizard, right-click the taskbar's network icon and select Connect to a network. If you select Wireless from the drop-down list, you'll see all the visible, hidden, and configured wireless networks on the machine. If a user attempts to connect to an unnamed (hidden) network, he or she will be prompted for an SSID name before authentication proceeds. Having to manually enter the SSID name every time you want to connect to a hidden network prevents broadcasting SSIDs from the client side when you're away from the network. You can automate this procedure by configuring Vista to connect automatically to hidden networks, although this approach requires broadcasting SSIDs. A better alternative it to use a semiautomatic approach: Configure the hidden network, deselect the option for automatically connecting to the network, but select the option to connect to the network even if it doesn't broadcast the SSID. To use this approach, select the Manage Wireless Networks option from Vista's Control Panel Network and Sharing Center applet, then open the wireless network's properties. This approach saves the network's SSID and authentication settings on the computer, but you still have to connect manually.

If you're wondering how Vista can discover hidden networks, then you should know that AP hardware actually hides SSIDs by sending a frame with the SSID set as NULL. Although XP and Windows Server 2003 can't display those networks to users, Vista can.

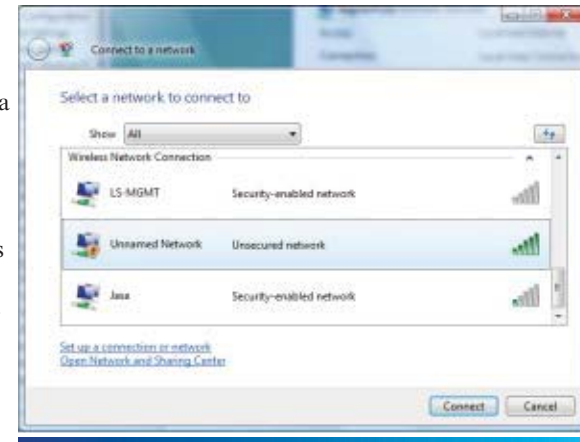

Figure 1: *Viewing hidden networks*

If a user tries to connect to an unsecured network, Vista notifies the user. A network is considered unsecured if it doesn't use an authentication and encryption protocol (or if it uses a weak protocol). A Vista client will never automatically connect to an unsecured network. You can use Group Policy to configure clients to prevent all unsecured connections. Automatic connections are possible only for secured networks that are configured with network profiles on the client side.

In Vista, creating and connecting to ad-hoc (without AP) networks is enhanced from both a security and a functionality standpoint. A major security feature for ad-hoc networks is implementation of the Wi-Fi Protected Access 2 (WPA2)-Personal security protocol. As Figure 2 shows, this protocol is the default authentication method in the wizard for creating ad-hoc networks. To access this wizard, start the Network and Sharing Center applet and select the Set up a connection or network option. Before Vista, Wi-Fi Protected Access (WPA) was available only on infrastructure wireless networks, and user-to-user networks were left with weak security methods such as static Wired Equivalent Privacy (WEP).

Another useful new feature for connecting Vista to wireless networks is Group Policy's Enterprise Single Sign-On service. This feature lets users authenticate to wireless networks and domain controllers (DCs) in a single logon procedure. First, the user is authenticated by using an 802.1x-enabled device (by using a certificate or a user name and password). If the logon is successful, the computer's Group Policy is applied, and credentials are passed to the domain logon procedure. Using the Enterprise Single Sign-On feature also lets you join a client to a domain by using only a wireless network, which isn't possible in XP. In XP, you have to connect the client to a physical network first, and join the client to the domain—then you can start to work on the wireless network.

Available Security Methods

Vista supports many security methods for authentication and encryption, as

Figure 2: *WPA2-Personal security protocol*

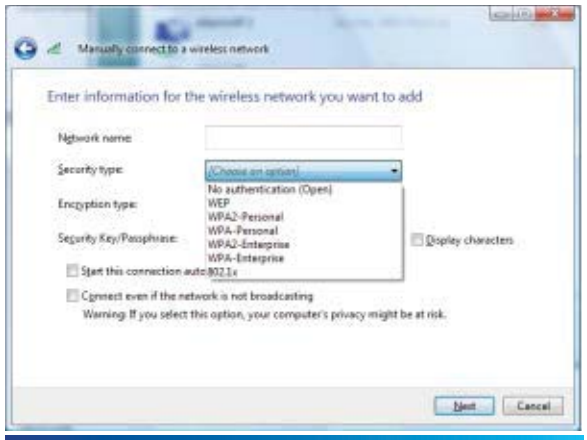

Figure 3: *Vista's security methods*

Figure 3 shows. WEP was the most commonly used security protocol for securing wireless networks in previous Windows versions. Although WEP is simple to implement, it's no longer considered a viable security method. WEP's main weakness is that it's based on a shared key for encryption of traffic (as well as for vector initialization). In addition, WEP uses an inferior encryption algorithm and has weak key management. These weaknesses make WEP an easily breakable solution that's no longer recommended.

The most commonly used security protocol in Vista is WPA. WPA has a better design, better key management, and a better encryption algorithm than WEP has. But WPA's major advantage over WEP is the use of Temporal Key Integrity Protocol (TKIP), which dynamically changes encryption keys as traffic goes between two hosts. Rather than WEP's cyclical redundancy check (CRC), WPA uses a better and more secure method for maintaining message integrity, called Message Authentication Code.

Vista offers two WPA configuration options: personal and enterprise. WPA-Personal is easier to configure because it uses a shared passphrase. This passphrase, which must be known (and configured) to the client and AP, acts as a base for implementing encryption. Although WPA-Personal is much more secure than WEP, sharing a passphrase can still pose a significant risk, so this implementation of WPA is recommended for small offices or home (ad-hoc) networks. WPA-Enterprise is a much more secure protocol, but it requires the implementation of 802.1x devices, the Remote Authentication Dial-In User Service (RADIUS) protocol, and an authentication server. WPA-Enterprise is intended for use in corporate environments. Both WPA-Personal and WPA-Enterprise also exist in version 2 (i.e., WPA2). The most important difference in version 2 is the implementation of the Advanced Encryption Standard (AES)-based algorithm, rather than WPA's RC4. Although WPA2 is recommended for optimal security, you might experience limitations if your AP or client hardware doesn't support it.

IEEE 802.1x authentication is designed for medium and large wireless LANs with authentication infrastructure consisting of RADIUS servers and account databases such as Active Directory (AD). This authentication method prevents a wireless client from joining a wireless network until it has performed a successful authentication. For authentication of clients, 802.1x uses EAP, with different methods such as those using user name and password credentials (Protected Extensible Authentication Protocol–Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MSCHAPv2) or a digital certificate and/or a smart card (Extensible Authentication Protocol–Transport Layer Security—EAP-TLS).

Using Group Policy to Manage Wireless Networks

Having a consistent policy for wireless connectivity in a corporate environment is important for maintaining a secure network. Using Group Policy is the easiest method for enforcing wireless and other policies. You can use Group Policy to block access to nearby wireless networks managed by different organizations, to disable the built-in support for wireless auto configuration, and to configure wireless clients to automatically connect to your organization's protected wireless networks.

In Windows 2003 and XP, you can use a Group Policy Object (GPO) to configure wireless settings. However, Windows 2003's GPO wireless options are limited to those available in XP. Vista greatly extends those capabilities, so the GPO now covers all the new features of wireless connections.

To use Group Policy for managing Vista wireless clients on a corporate level, you must first extend Windows 2003's AD schema with the proper attributes. The Microsoft article "Active Directory Schema Extensions for Windows Vista Wireless and Wired Group Policy Enhancements" (www.microsoft.com/technet/network/wifi/vista_ad_ext.mspx) includes detailed instructions for this procedure, as well as the required script. After you extend the AD schema, you can use Vista's Group Policy Management Console (GPMC—connected to the corporate forest) to configure wireless policies. Create a new GPO, then navigate to Computer Configuration, Windows Settings, Security Settings, Wireless Network (IEEE 802.11) Policies. Because Vista has a new set of wireless options, you must create separate policies for XP and Vista. Fortunately, you don't have to create a separate GPO for each OS and deal with WMI. You can simply right-click the GPO Wireless Network Policies item and create a new XP or Vista policy. If both types of wireless policies are configured, XP wireless clients will use only their own policy settings, and Vista wireless clients will use only their own policy settings. If no Vista policy settings exist, Vista wireless clients will use the XP settings, because they're a subset of the settings available for Vista. Note that wireless policies intended for Vista, created from Vista's GPMC and linked somewhere in domain, aren't visible from Windows 2003's GPMC (unlike XP policies). However, this doesn't mean that the policies won't be

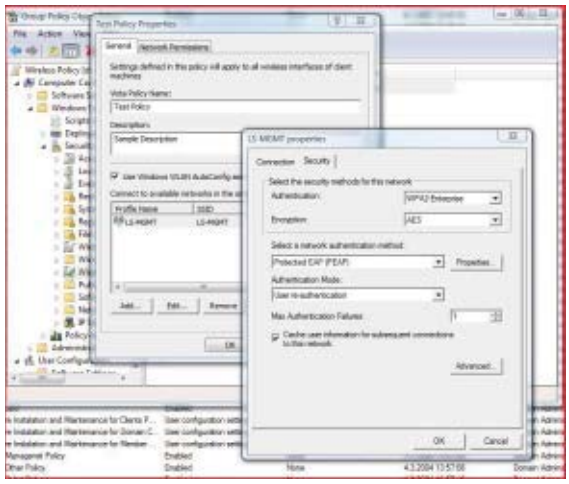

Figure 4: *Configuring wireless policies*

applied.

Wireless policies have many configuration options, such as preventing users from connecting to ad-hoc networks, preventing users from creating new wireless profiles, and enforcing only preconfigured wireless profiles. By using these options in Group Policy, administrators can create wireless profiles for some or all users that contain information about the SSID, authentication and encryption methods, and some advanced 802.1x options. For example, if you want to preconfigure a wireless network profile for a client so that he doesn't have to enter any settings, open a new policy window, select the General tab, click Add, and select the network type (infrastructure or ad-hoc). Then, enter all the data for the desired wireless network in the new profile properties window that opens (which Figure 4 shows an example of). If you want to restrict users to connect only to networks that you explicitly specify, select the Network Permissions tab rather than the General tab.

Using Group Policy is the only method for configuring Vista's Enterprise Single Sign-On feature. Enterprise Single Sign-On options in Group Policy let you configure when 802.1x authentication will occur in relation to user logon, as well as let you integrate user logon and 802.1x authentication credentials on the DC. You can choose between performing wireless authentication immediately before or after user logon, and you can specify the number of seconds of delay for connectivity before the process begins. You can also configure options to prompt the user to fill in additional fields if necessary, and you can specify whether your wireless networks will use a different Virtual LAN (VLAN) for computer and user authentication. To configure these options, open a new policy window, select the General tab, click Add, and select Infrastructure. In the new profile properties window that opens, select the Security tab and click Advanced.

If you're using WPA2-Enterprise authentication, Group Policy offers a set of options for configuring the caching of 802.1x authentication results, as Figure 5 shows. In the Fast Roaming section, you can configure Pairwise Master Key (PMK) caching and preauthentication options. Wireless clients and wireless APs can both cache the results of 802.1x authentications. Caching those results makes subsequent access much faster when a wireless client roams back to a wireless AP to which the client already authenticated. You can configure a maximum time to keep an entry in the PMK cache and the maximum number of entries. With preauthentication, a wireless client can perform an 802.1x authentication with other wireless APs in its range while it's still connected to its current wireless AP. You can also configure the maximum number of times to attempt preauthentication with a wireless AP.

Wireless Networks and NAP

Network Access Protection (NAP), which is Windows Server 2008's and Vista's new feature for controlling network access (from the client health aspect), can also be applied to wireless networks. Vista can declare its health state while trying to connect to 802.1x-enabled wireless networks. For NAP to work on a wireless network, the current domain environment must include Server 2008 Network Policy Server (NPS). On the client side, Vista must be configured with the proper enforcement agent for 802.1x (i.e., the EAP Quarantine Enforcement Client). To configure this enforcement agent, open the NAP Client Configuration console (napclcfg.msc) and go to the Enforcement Agents node. Start the Services applet from the Control Panel's Administrative Tools, and configure the Network Access Protection service to start automatically.

When a client that doesn't comply with company security requirements (e.g., doesn't have all updates installed) tries to connect to the corporate wireless network, NAP will deny access and will place the client in quarantine (on a separate VLAN). The client will be able to access only remediation servers (e.g., Windows Server Update Services—WSUS) that will provide the necessary updates to make the client compliant. For more information about NAP, including configuring NAP with 802.1x (which is beyond the scope of

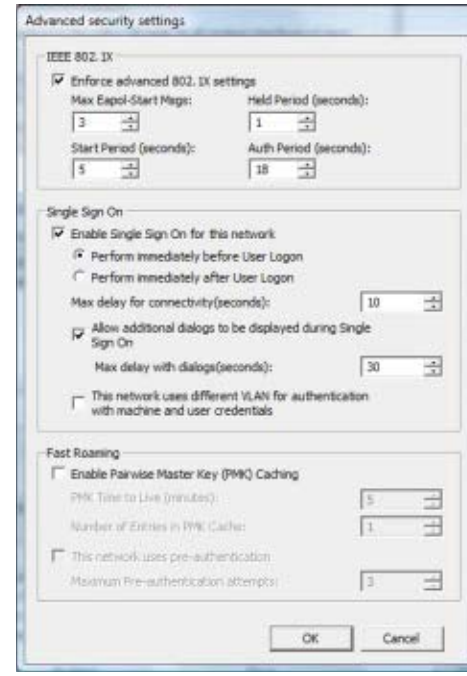

Figure 5: *Configuring advanced security settings for a wireless GPO*

this article), go to technet.microsoft.com/en-us/network/bb545879.aspx.

Unplug Safely

Vista's new wireless features can help enhance wireless security in both home and corporate environments. Implementing WPA2 in ad-hoc networks can improve home network security. For corporate implementations, Vista can work with the latest security technologies such as NAP to boost wireless security.

InstantDoc ID 97336

Damir Dizdarevic *(ddamir@logosoft.ba) is the manager of the Learning Center at Logosoft in Sarajevo, Bosnia. An MCSE, MCTS, MCITP, and MCT, he specializes in Windows Server security and has published more than 350 articles in IT magazines.*

Vista and Server 2008 Malware Protection Gems

USE DEP AND ASLR TO PROTECT YOURSELF AGAINST BUFFER-OVERRUN-BASED ATTACKS

By Jan De Clercq

Attacks based on buffer overruns (aka buffer overflows) have been a problem for a long time and are still considered one of the computer industry's most important security problems. The first buffer-overrun–based attack distributed via the Internet, the Morris worm, did a lot of harm in 1988. The sad thing is that the creators of the Morris worm didn't write the worm to cause harm but rather as an experiment for measuring the size of the Internet. The Morris worm exploited weak passwords and known vulnerabilities in UNIX programs such as sendmail and Finger. Two recent well-known attacks that involved exploiting buffer overruns, the Code Red and SQL Slammer worms, exposed many Internet-connected systems to attackers' control. In 2001, the Code Red worm exploited a buffer-overrun vulnerability in Microsoft Internet Information Services (IIS) 5.0 (the IIS version that is bundled with Windows 2000), and in 2003, the SQL Slammer worm used a buffer-overrun vulnerability to compromise machines running Microsoft SQL Server 2000.

You can defend against buffer-overrun–based attacks by using defenses that Microsoft includes in Windows Vista and Windows Server 2008: Data Execution Prevention (DEP) and Address Space Layout Randomization (ASLR). (At the time of this writing, Microsoft was about to release Vista SP1 and had released Windows Server 2008 RC0.) I'll explain why these defenses are important and how you can configure them and observe their behavior.

Understanding Buffer Overruns

Before going into more detail on the Vista and Server 2008 buffer-overrun defenses, it might be worthwhile to look at how a buffer overrun works and how it can harm your systems and data.

A buffer overrun occurs when a malicious or badly engineered program stores data beyond the boundaries of a fixed-length buffer in computer memory. The result is that the extra "overflowing" data overwrites adjacent memory locations. The data that's overwritten can include other buffers, variables, and program logic and may cause a process to crash or produce incorrect results. An even bigger threat is that the injected data often includes executable code that the program under attack is then lured to execute. This executable code often contains the real payload of a buffer-overrun–based attack. It's used to steal or delete data, create: Denial of Service (DoS)–based service outages, trigger privilege elevations, or spread malware to other systems.

Figure 1 gives a simple example of a buffer overrun. A program has defined two variables that are stored in adjacent memory locations. The first variable is an eight-byte-long string called X; the second, a two-byte integer called Y. Initially, X contains nothing but zero bytes, and Y contains the number 30. Imagine that a user (whether unintentionally or maliciously) inputs a character string OVERFLOW to this program. The program then attempts to store this character string in X's memory location followed by a 0 value to mark the end of the string. The program logic doesn't check the length of the string and partially overwrites the value of Y. The result is that, although the programmer didn't intend to change the value of Y when variable X receives input, variable Y's original value 30 is now replaced by the number that's part of the character string that was injected into the variable X memory location.

Developers can prevent buffer overruns by including sufficient boundary checks in their program code and by leveraging compilers or runtime services that perform boundary checks. Boundary checks ensure that input data are of the right length. Although boundary checking and enforcement have become best practices for developers, plenty of legacy code doesn't include boundary checks. Also, coding best practices are worthless if some programmers don't follow them.

These reasons explain why many hardware, application, and OS software vendors including Microsoft have developed proactive defenses that

Figure 1: *Simple buffer overflow example*

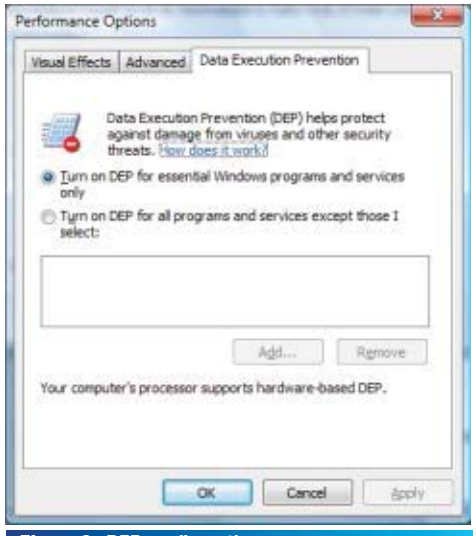

Figure 2: *DEP configuration*

attempt to stop buffer-overrun attacks in badly engineered code. Let's look at two Microsoft implementations of proactive buffer overrun defenses: DEP and ASLR.

Data Execution Protection

As I mentioned above, buffer-overrun–based attacks often write executable malicious code to another program's memory buffers and then trick the program into executing the malicious payload. You can tackle the execution of maliciously injected code by using a technique that Microsoft refers to as Data Execution Protection (DEP). DEP lets Windows mark memory locations that should only contain data as non-executable (NX). When an application attempts to execute code from NX-marked memory locations, Windows' DEP logic will block the application from doing so.

A negative side effect of the buffer-overrun protection offered by DEP is that the blocked application will typically halt. In other words, even though DEP stops malware from executing its malicious payload, this situation creates a new opportunity for malware to launch DoS attacks.

Microsoft includes DEP support not only in Vista and Server 2008, but also in Windows XP SP2, Windows Server 2003 SP1, Windows 2003 R2. Microsoft DEP implementation comes in two flavors: hardware-enforced DEP and software-enforced DEP.

Hardware-enforced DEP. Hardware-enforced DEP leverages a processor feature that AMD refers to as the no-execute page-protection (NX) feature and that Intel refers to as the Execute Disable Bit (XD) feature. At the time of writing, AMD supported NX only on its 64-bit processors, and Intel supported XD only on the Itanium and EM64T 64-bit processors and a small number of 32-bit Prescott processors. Microsoft is not the only OS vendor that leverages the NX and XD processor features for stopping buffer overruns: NX- and XD-enabled software is also available in other OSs such as Linux and UNIX BSD (see en.wikipedia.org/wiki/Nx-bit for more information).

Software-enforced DEP. Software-enforced DEP lets Microsoft provide DEP on 32-bit processor systems not equipped with an NX- or XD-compatible processor. In this software workaround, the processor-level NX- or XD-bit functionality is provided by a set of special pointers that the Windows OS automatically adds to data objects stored in the system memory.

You can easily check whether your system supports hardware- or software-enforced DEP by checking the DEP configuration settings. You can access these settings using the Advanced Settings option in the System Control Panel applet and navigating to the Advanced and Performance Settings options.

At the bottom of the DEP configuration settings screen, there's a reference to the type of DEP your system supports. Figure 2 shows the DEP configuration settings on a Vista system. (I'll explain the other configuration options later in this section) The bottom line reads, "Your computer's processor supports hardware-based DEP."

If your system supports software-enforced DEP (meaning that your machine doesn't have the NX- or XD-compatible processor), you'll see "Your computer's processor does not support hardware-based DEP. However, Windows can use DEP software to help prevent some types of attacks."

An alternative way to check whether your system supports hardware- or software-enforced DEP is by using Windows Management Instrumentation (WMI) commands. The procedure is outlined in the Microsoft article at support.microsoft.com/kb/912923.

On XP SP2, Windows 2003 SP1, and later Microsoft OSs, DEP is enabled by default. However, DEP doesn't always protect all programs running on your system. The exact list of programs that are protected by DEP is defined by DEP's protection level. DEP supports two protection levels:

- Level 1—The first level protects only the Windows system code and executables and doesn't offer DEP

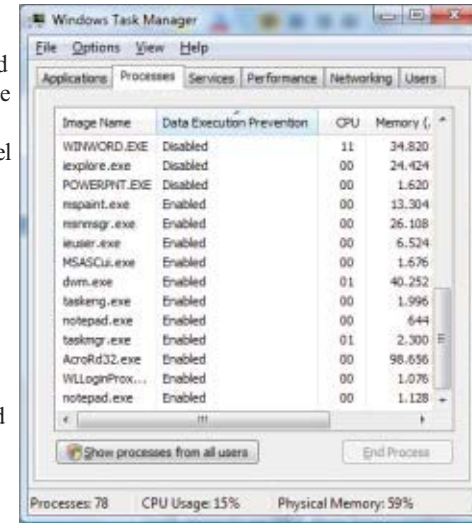

Figure 3: *Checking DEP status of a process from the task Manager*

/NoExecute= values	Meaning
AlwaysOn	DEP always turned on for all services and applications – grays out the DEP configuration screen (see figure 2) in the system properties
AlwaysOff	Completely turns off DEP
OptIn	Turns DEP on and sets it to protection level I
OptOut	Turns DEP on and sets it to protection level 2

Table 1: *Boot.ini NoExecute= values and Their Meaning*

protection for additional Microsoft or third-party applications that run on your system.

- Level 2—The second level protects all executable code that runs on your system; it offers DEP protection for both Windows system code and the Microsoft or third-party applications that run on your system.

By default, XP SP2 and Vista run DEP at protection level 1; Windows 2003 SP1 and Server 2008 run DEP at protection level 2.

Administrators can configure the DEP protection levels from the DEP configuration screen, which you can see in Figure 2. In this example (which shows the default DEP configuration settings on a Vista system), DEP is enabled for essential Windows programs and services only—DEP protection level 1. You can use the other radio button Turn on DEP for all programs and services except those I select to switch to DEP protection level 2, which is the default setting on Windows 2003 SP1 and Server 2008.

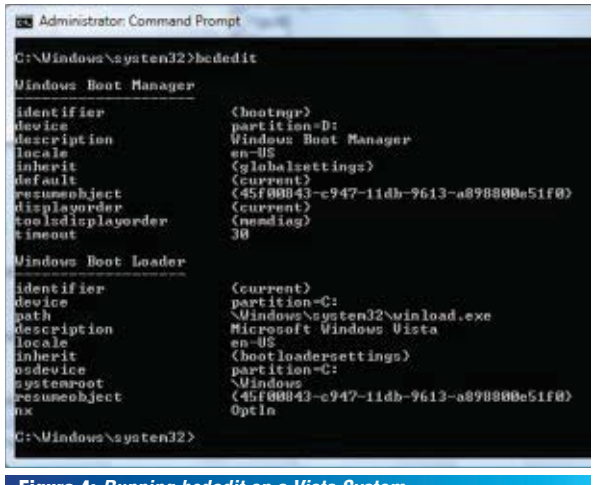

Figure 4: *Running bcdedit on a Vista System*

Figure 2 also shows that protection level 2 lets you exempt certain applications from DEP protection. This ability to exempt apps is important because some legacy applications don't run properly when DEP is enabled—for example, at the time of writing, Microsoft Word was still automatically exempted from DEP. Before switching your DEP protection to level 2, you must run an application compatibility test to ensure that all applications run properly when DEP is enabled. To exempt one of your applications from DEP, you can add the application's executable to the excluded list in the DEP configuration screen using the Add… button.

You can easily check whether a given application is protected by DEP by checking the DEP column of the application's process in the Windows Task Manager, which Figure 3 shows. If you don't see the DEP column on your system, you can add it using the Task Manager's View\Select Columns… option.

Another way to exempt one of your applications from DEP is to create a software fix to distribute to your systems that automatically disables DEP for a given application on those systems. Microsoft refers to such a software fix as a DisableNX shim. To create this software fix, see the Microsoft Application Compatibility Toolkit (ACT), which also includes a tool called Compatibility Administrator that can help (technet.microsoft.com/en-us/windowsvista/aa905078.aspx). Application developers can also do the opposite—directly enable their applications for DEP support in their application binaries. To do so, they use the /NXCompat compilation switch.

One important final note is that when DEP is running in protection level 2, your system will run a bit slower because of all the extra DEP checks that are carried out on the processor and system memory level. That's why for test systems that aren't exposed to the Internet, for example, you can consider turning off DEP protection completely. The only way to turn off DEP completely on a given system is to specify the "/NoExecute=AlwaysOff" switch in the system's boot.ini file. Note that you can also use the same boot.ini /NoExecute= switch with other values to turn DEP on and to set the DEP protection level. Table 1 shows all the /NoExecute values.

The boot.ini file is available only on XP and Windows 2003, and you can edit it using Notepad or going to the Startup and Recovery section in System properties. On Vista and Server 2008, the boot.ini has been replaced by the Boot Configuration Data (BCD) file. To edit the BCD file, Microsoft provides a command-line utility called bcdedit.exe. When you run bcdedit without switches, it shows your current boot configuration. Figure 4 shows the result of running bcdedit on a Vista system. Note the last line that holds the nx configuration: To change the NX configuration to alwaysoff, for example, you would run the following bcdedit command:

```
bcdedit /set nx alwaysoff
```

The values specified in Table 1 for the boot.ini /NoExecute= switch are also available for the BCD nx option.

For more information about Microsoft DEP and how to configure it, consult the Microsoft article at support.microsoft.com/kb/875352/en-us.

Address Space Layout Randomization

Another technique often used by buffer-overrun–based malware is to inject a system memory path that points to the location of an important system DLL into another program's buffer. The malware then tricks the

program into calling that particular system file to let the malware leverage the system DLL's services without being detected.

This type of buffer-overrun attack is relatively easy to carry out if the OS always loads certain system DLLs on the exact same memory location. On XP, for example, the memory locations of system DLLs are always identical— they vary only slightly depending on the service pack status of the system. The new Vista and Server 2008 Address Space Layout Randomization (ASLR) feature makes it harder for malware to leverage a system DLL's services by randomizing the DLLs' memory location. Unlike DEP, ASLR isn't available on earlier Windows versions.

Each time a Vista and Server 2008 system reboots, ASLR randomly assigns system code (basically system DLLs and executables) to different memory locations. This means that the system code's entry points (the addresses the malware would use to call on the service of a particular piece of system code) are in unpredictable locations. In Vista and Server 2008, a DLL or EXE can be loaded into any of 256 locations. This means that an attacker has a 1/256 chance of getting the address right. As such, ASLR also makes it harder for hackers to write repeatable code such as worms that target identical system resources on many different systems.

You can observe the effect of ASLR by using the SysInternals Process Explorer tool, which you can download at www.microsoft.com/technet/sysinternals/utilities/processexplorer.mspx. To use the tool, start Process Explorer and ensure that you have selected the Show Lower Pane option in the View menu.

Then select the explorer.exe process in the upper pane and check the base address of the ntdll.dll in the base column in the lower pane. (If you don't see the base column you can add it by using the View / Select Columns… menu option—the Base column can be added from the DLL tab.)

Write down the base address, then reboot your system. On an XP system, the base address for ntdll.dll remains identical after a system reboot (XP doesn't support ASLR). On a Vista system, the base address is different after a system reboot (Vista supports ASLR).

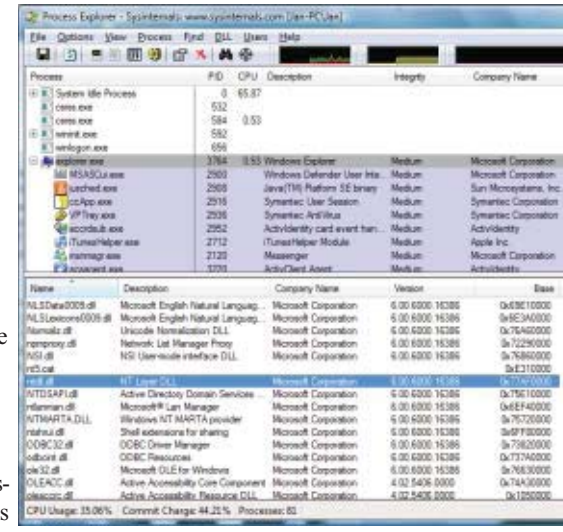

Figure 5: *Observing the effect of ASLR with SysInternals Process Explorer*

Figure 5 shows the Process Explorer interface and the base address for the ntdll.dll DLL. Table 2 shows the base addresses I found for the ntdll.dll and user32.dll DLLs when running Process Explorer on an XP SP2 system and on a Vista system.

You can leverage ASLR not only for randomizing the memory locations of Windows system files but also for randomizing the memory locations of executables and DLLs of any application that runs on Vista or Server 2008. To do so, application developers must compile their code with the /dynamicbase linker option. Microsoft Visual Studio supports this option from Visual Studio 2005 SP 1 and later.

Like DEP, ASLR is not a Microsoft-only invention and implementation. ASLR was implemented long before Vista and Server 2008, on platforms such as Linux and UNIX. Also certain Host Intrusion Detection System (HIDS) solutions have been supporting ASLR on legacy Windows platforms long before the native Windows support. A good analysis of the Microsoft ASLR implementation in Vista is offered in the Symantec research paper at www.symantec.com/avcenter/reference/Address_Space_Layout_Randomization.pdf. Unlike with DEP, Microsoft doesn't offer ASLR-specific configuration settings for fine-tuning the use of ASLR.

Important Proactive Defenses

DEP and ASLR each use a slightly different proactive defense approach as a buffer-overrun defense. Where ASLR makes it more difficult for malware to find the right code, DEP makes it more difficult for malware to execute the code once the target code is found. You can leverage both techniques at the same time and they can also be leveraged in virtual computing environments such as Microsoft Virtual PC or VMware products.

From an application-support point of view, you should remember that you must test your applications for DEP compatibility prior to deploying them on a DEP-enabled Windows platform. DEP can cause certain applications to stop working properly or even halt.

DLL	Windows XP SP2 base address No ASLR	Windows XP SP2 base address (after reboot) No ASLR	Windows Vista base address With ASLR	Windows Vista base address (after reboot) With ASLR
Ntdll.dll	0x7C900000	0x7C900000	0x77AF0000	0x776B0000
User32.dll	0x7E410000	0x7E410000	0x76880000	0x76520000

Table 2: *Effect of ASLR on DLL Base Addresses*

Finally, it's important to understand that DEP and ASLR aren't a panacea for the buffer-overrun problem. Both techniques certainly make it much more difficult for malware to leverage buffer overruns. ASLR, for example, doesn't make it impossible for malware to find system code, but it certainly makes the process of finding it much more challenging. But in many cases ASLR and DEP will also effectively stop buffer-overrun–based attacks.

InstantDoc ID 98005

Jan De Clercq *(jan.declercq@hp.com) is a member of HP's Security Office and focuses on identity management and security in Microsoft products. He is coauthor of* Microsoft Windows Security Fundamentals *(Digital Press).*

6 New Security Features in IIS 7.0

HOW THEY CAN HELP YOU GET CONTROL OVER YOUR WEB SERVER AND REDUCE YOUR ATTACK SURFACE

By Derek Hatchard

When you host a Web server, you put a part of your organization on display and open it up to the poking and prodding of the anonymous masses. Remotely exploitable flaws in the Web server platform can be disastrous. Case in point: Microsoft Internet Information Services (IIS) 5.0 left a trail of lost productivity and revenue.

However, Microsoft redesigned IIS with security as a top priority. The result was IIS 6.0, which is widely held as the most secure commercial Web server on the market (as indicated by the low number of Secunia advisories about it—three—see secunia.com/product/1438).

IIS 7.0 builds on the secure design of IIS 6.0 and has been modularized so that individual features can be removed entirely, thus reducing the overall attack surface of your Web server. Application pools, introduced in IIS 6.0 as a way to isolate applications from each other (and from the Web server process), are now more effectively sandboxed. New delegation features let site owners manage their sites without elevated privileges. Request filtering (aka URLscan) is now built into the server. And administrators can define rules right in IIS 7.0 that control which users have access to which URLs.

These features are among the security-related enhancements in IIS 7.0. They're worth a closer look, and they might even change the way you think about managing and configuring Web sites.

Application Sandboxing

Consider a market research company hosting surveys or other low volume sites for competing companies on the same box. Or consider a server that hosts a payroll application used by a small number of users and a homegrown portal used companywide. In both cases it's crucial that these applications running on the same servers be isolated from each other.

Web applications run in worker processes. Application pools map Web applications to worker processes. A specific worker process is used only to run applications that are part of the same application pool. In IIS 6.0 and IIS 7.0, the worker process is w3wp.exe.

In IIS 6.0, new Web sites and applications are put into the same application pool. This default application pool runs under the NetworkService account. As an administrator, you can create new application pools manually and assign Web apps to those pools. By default, those application pools will also run under the NetworkService account, which can lead to an undesirable runtime scenario as all Web applications run with the same permissions. An application in app pool A can read the configuration of app pool B and even access the content files of applications assigned to app pool B. Although it's easy enough to create new app pools and to configure custom accounts for each, managing those accounts over time is cumbersome.

With IIS 7.0, a new application pool is created automatically for each Web site. By default, that application pool is configured to run as the NetworkService account. But when the worker process is created, IIS 7.0 injects a special SID unique to the app pool into the NetworkService security token. IIS 7.0 also creates a configuration file for the worker process and sets the file's ACL to allow access only to the unique SID for the app pool. The result is that an application pool's configuration can't be read by other application pools.

As an additional precaution, you can change the ACLs on content files to provide access to the unique app pool SID instead of NetworkService. This will prevent an application in app pool A from reading the content files of an application in app pool B.

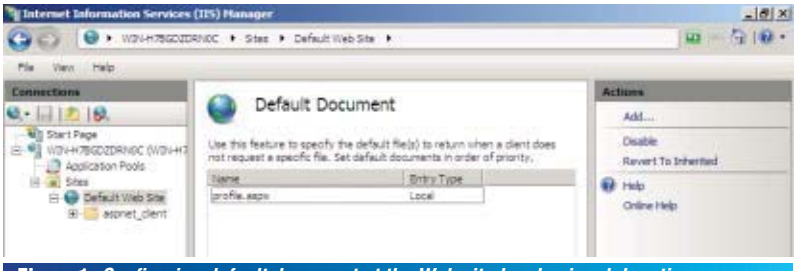

Figure 1: *Configuring default document at the Web site level using delegation*

IUSR and IIS_IUSRS

Tangentially related to process identity is the question of which identity the server uses for anonymous requests. Previous versions of IIS relied on a local account, IUSR_servername, as the identity for anonymous users. IIS 7.0 uses a new built-in account called IUSR. You can't log in locally with the IUSR account, so it doesn't have a password (which means there are no risks due to attackers guessing the password). The IUSR account always has the same SID so ACLs are transferrable between Windows Server 2008 machines (as well as Windows Vista machines). And if the IUSR account isn't appropriate for your scenario (e.g., if anonymous requests require authenticated network access), you can turn off the anonymous user account and IIS 7.0 will use the worker process identity for anonymous requests.

Also new is the built-in IIS_IUSRS group. This group replaces the IIS_WPG group. In IIS 6.0, the IIS_WPG group provides the minimum rights needed to run a worker process, and you must manually add an account to this group to provide a custom identity for a worker process. The IIS_IUSRS group provides a similar role for IIS 7.0, but you don't explicitly add accounts to this group. Instead, IIS 7.0 automatically enrolls accounts in IIS_IUSRS when they're assigned as the identity for an application pool. And as with the IUSR account, the IIS_IUSRS group is built-in, so it always has the same name and SID on all Server 2008 installations, making ACLs and other configurations completely portable between Server 2008 machines (and Vista machines).

Feature Delegation

Not every Web server setting really needs to be protected by admin rights. Some settings are simple application-level decisions that can be made by developers or product managers. For example, in IIS 6.0 you need admin rights to change the default document for a Web application. But normally is there really any reason that the ability to change default.aspx to profile.aspx should require administrator rights?

In IIS 7.0, configuration decisions can now be delegated to site or application owners. IIS 7.0 uses a new XML-based configuration system inspired by ASP.NET. At the site and application level, both IIS 7.0 and ASP.NET configuration settings are found in the same web.config files.

Delegated settings such as the default document can be changed at the Web site level or application level by editing the web.config file directly or using the IIS Manager GUI, as Figure 1 shows, which updates the web.config for you. In the web.config file, the system.webServer section contains the IIS 7.0 configuration settings, which Figure 2 shows.

The sections that are valid within <system.webServer> are defined in a special configuration file called applicationHost.config. In applicationHost.config, each section has a default delegation mode. In the example in Figure 3, the default document and directory browsing settings can be overridden but not the asp, caching, or cgi sections because they're in Deny for that setting.

But what if there is a good reason to prevent a Web site owner from changing the default document? No problem: IIS 7.0 lets you lock configuration elements so they can't be set or overridden in web.config files. In the case of the default document, you can globally change the default override mode to Deny or you can explicitly set the override mode to Deny for specific locations (using location tags). The IIS team recommends asserting these kinds of changes in location tags, as Listing 1 shows. Feature delegation can be a great boon to a busy administrator because it safely empowers Web site and application owners to configure aspects of the Web server that affect only their sites and applications.

Administration Delegation

Many admins find it expedient to just give out admin access to whoever needs to apply a change to a site or application. This, of course, is a tremendous security risk. Unfortunately, the choice has been difficult: either liberally assign admin rights or impede updates by becoming the single point of administration. Now, though, feature delegation in IIS 7.0 solves much of the problem by allowing site owners to put configuration in web.config files. With IIS 7.0, server admins can also grant administration rights for a specific Web site or application to one or more users without elevating user privileges.

With IIS Manager, which Figure 4 shows, users can connect to an IIS 7.0 server using Windows credentials or credentials specific to IIS Manager. The beauty of credentials specific to IIS Manager is that you provide a very specific and limited set of rights to a user: IIS Web site administration rights, which are useless outside of IIS Manager.

For remote use, a standalone version of IIS Manager is available for Windows Vista, 2003, and XP. Before you can connect remotely with IIS Manager, remote management must be explicitly enabled on the Web Server by doing the following:

1. Install the Web Management Service (WMSVC)
2. Enable remote management via IIS Manager on the Web server (or via the registry)

```
<system.webServer>
<defaultDocument>
<files>
<clear />
<add value="profile.aspx" />
</files>
</defaultDocument>
<directoryBrowse enabled="true" />
</system.webServer>
```

Figure 2: *IIS 7.0 Configuration Settings*

```
<sectionGroup name="system.webServer">
<section name="asp" overrideModeDefault="Deny" />
<section name="caching" overrideModeDefault="Deny" />
<section name="cgi" overrideModeDefault="Deny" />
<section name="defaultDocument" overrideModeDefault="Allow" />
<section name="directoryBrowse" overrideModeDefault="Allow" />
```

Figure 3: *Default Settings in Override Mode Deny and Override Mode Allow*

```
<location path="Default Web
Site" overrideMode="Deny">
<system.webServer>
<defaultDocument>
<files>
<clear />
<add value="default.aspx" />
</files>
</defaultDocument>
</system.webServer>
</location>
```

Listing 1: *Using Location Tags to Set Override Mode to Deny*

```
<location path="Reporting">
<system.webServer>
<security>
<authorization>
<remove users="*" roles="" verbs="" />
<add accessType="Allow" roles="Managers" />
<add accessType="Deny" users="*" />
</authorization>
</security>
</system.webServer>
</location>
```

Listing 2: *URL Authorization Rules Syntax*

3. Start the Web Management Service

Firewall rules or remote access policies can make it difficult to use remote management tools. For this reason, IIS Manager works over HTTPS, so it's both secure and firewall-friendly. By default, the Web Management Service uses a self-signed certificate and listens on port 8172.

Microsoft offers IIS 7.0 Manager for remote management at www.iis.net/go/1524. For additional resources (including detailed configuration instructions), search for IIS 7.0 remote administration at iis.net. You can also find more information about the new IIS features at this Microsoft site.

Built-In Request Filtering

If you've administered IIS servers, you're probably already familiar with UrlScan, a downloadable tool for IIS 4.0 and higher that restricts the types of requests that IIS will service. The intent behind request filtering is to protect your Web server from potentially malicious requests.

In IIS 7.0, UrlScan has been enhanced and bundled with the Web server in the Request Filtering Module. The Request Filtering Module rejects requests based on configurable criteria. For example, the module can reject double-encoded requests or requests of unusual size (such as large POST pay-loads or URLs that are too long). The Request Filtering Module can also reject requests for file types, paths, or HTTP verbs that your site doesn't support.

With IIS 7.0, request filtering configuration can be delegated, allowing site admins to define their own request filtering rules in web.config files, which wasn't possible with UrlScan and IIS 6.0. For more information about request filtering in IIS 7.0, see the Security ProVIP article "Unleash the Power of IIS 7.0's Security Features," InstantDoc ID 96999.

Figure 4: *IIS Manager*

URL Authorization

Web applications often have restricted areas to which only certain users have access. Only a manager, for example, is allowed to access performance reviews in an HR system. These restricted pages are commonly grouped together into directories with names like Administration, Reporting, or Moderation. Properly securing these sections to prevent unauthorized access has been cumbersome at best with previous versions of IIS. Even with the URL authorization feature built into ASP.NET, you still have to deal with non-ASP.NET content such as PDF or Excel files that need to be protected. And ASP.NET URL authorization rules are managed by editing XML, which can be tedious.

In IIS 7.0, ASP.NET URL authorization is still available, but in addition, a URL authorization feature is provided by the Web server itself. Now access to all content types (e.g., static, PHP, ASP) can be controlled based on user, group, and URL. For example, you can easily restrict access to anything under the Reporting path to only those users belonging to the Managers group--without touching the file ACLs. Figure 5 shows URL authorization rules configuration in IIS Manager.

URL authorization rules are persisted in the system.webServer section of web.config files with a slightly different syntax than ASP.NET authorization rules, as Listing 2 shows.

Since the authorization rules are contained entirely in your configuration files (local web.config), they are easily transferred between applications and servers. And the URL authorization in IIS 7.0 works with Windows users and groups as well as ASP.NET users and roles.

Building on Solid Security

IIS 7.0 builds on the solid security legacy of IIS 6.0. IIS 7.0 retains the core architecture of IIS 6.0 with the

Figure 5: *URL Authorization rules configuration in IIS Manager*

app pool / worker process isolation model that has proven to be very effective. On top of that secure foundation, Microsoft has introduced a number of important security improvements. Although the new modular architecture changes receive a lot of the attention when discussing IIS 7.0 security, features such as automatic application sandboxing, feature delegation, and URL authorization make it easier than ever to ensure your Web server is secure during day-to-day operations.

InstantDoc ID 98393

Derek Hatchard *is a Web entrepreneur, consultant, and trainer/mentor. He is also a Microsoft Regional Director, Microsoft MVP, author, and speaker. He blogs at derekh.com and ardentdev.com, and he co-hosts a developer podcast at devcasting.com.*

Windows Server 2008 Password Policies

THE NEW SERVER OS RESOLVES EARLIER PASSWORD POLICY LIMITATIONS

By Jan De Clercq

One of Windows' most important security policies that every Windows administrator is certainly familiar with is the password policies. These policies let you enforce password quality requirements (e.g., minimum password length, maximum password age) for the passwords of local or domain user accounts. As you might know, Windows Server 2003 and Windows 2000 Server password policies have some important limitations. In this article I explain these limitations and discuss how Windows Server 2008—Microsoft's upcoming Server OS—resolves them. I also explain how you can configure and use Server 2008's password policies. At press time, Microsoft had released Server 2008 Release Candidate 0 (RC0) and was planning to launch the Server OS on February 27, 2008.

A Flexible Solution for a Serious Problem

A serious limitation of the password policies in Windows 2003 and Win2K is that administrators can define only one password policy that applies to all user accounts in a domain. You can define this global domain password policy from the Default Domain Policy Group Policy Object's (GPO's) Password Policy settings or from any other GPO that's linked to the Active Directory (AD) domain object. To access the Password Policy configuration interface, go to the \Computer Configuration\Windows Settings\Security Settings\Account Policies GPO container. Even though you can define different password policies in the GPOs and link them to AD organizational units (OUs) or computer accounts, these password policies don't apply to domain accounts—instead, they apply to the local accounts that are defined in the security databases of the computer accounts to which the GPOs apply.

Organizations typically want to impose different password quality requirements for certain categories of domain accounts. A classic example is having a different password policy for administrator accounts and regular user accounts. The security rationale is simple: Administrator accounts have more powers (permissions and rights) than plain user accounts, so you might want a higher quality authentication process for administrators than for regular users. Another way to provide stronger authentication is to enforce the use of smart card logon for administrator accounts.

Windows 2003 and Win2K provide two workarounds for organizations that want to define different passwords policies in a single domain, although both workarounds are difficult to implement. One workaround is to deploy separate domains for each of the account categories that you want to define a special password for. The other workaround is to develop a special "password filtering" DLL that you then deploy to all your domain controllers (DCs). The second solution is rarely used because it's even more complex and time consuming than the first solution.

Server 2008 comes to the rescue by introducing fine-grained password policies that let administrators define different password policies for different domain account categories in a

single domain. This new fine-grained password policy functionality can be applied only to domain accounts—not to local accounts.

Server 2008 introduces the same functionality for the account lockout policies that in earlier Windows Server versions were crippled by the same limitation (i.e., you could define only a single account lockout policy for all domain accounts). Account lockout policies ensure that user accounts automatically become unusable after a user enters a certain number of incorrect passwords. The administrator must define a bad password threshold to configure the account lockout policy.

Configuring Fine-Grained Password Policies

Configuring Server 2008's fine-grained password policies is entirely different from defining the classic domain account or local account password policy in earlier Windows versions (which I described previously). You can't use GPO settings to configure fine-grained password policies, because Microsoft uses a different (non–GPO-based) mechanism to store and enforce these policies.

Server 2008's fine-grained password policies are stored in a new AD container called the AD Password Settings Container, which is located in the System container of the AD domain naming context. To define a new fine-grained password policy, you must create a new AD object of the msDS-PasswordSettings object class in this container. Objects of this class are referred to as Password Settings objects (PSOs) in the Microsoft documentation. By default, only members of the Domain Admins group can create PSOs, because only members of this group have the AD Create Child and Delete Child permissions on the Password Settings Container. (I discuss the tools you can use to create and configure PSOs in a later section.)

To apply the PSOs you created, you must link the PSO to an AD user or group object. To do so, you don't need permissions to the AD object itself; you simply need Write permissions on the PSO. By default, only members of the Domain Admins group have this permission. Therefore, only members of the Domain Admins group can link a PSO to a group or user—although you can obviously delegate these permissions to other administrators.

Table 1 summarizes the attributes that are linked to Server 2008 PSOs. Note that a PSO can store not only password policy settings but also account lockout policy settings. Remember that Server 2008 supports both fine-grained password and account lockout policies. Two important PSO attributes are the msDS-PSOAppliesTo and msDS-PasswordSettingsPrecedence attributes.

The msDS-PSOAppliesTo PSO attribute is a multi-valued attribute

Attribute Name	Required?	Description	Example Value
cn	Yes	Common name	MyPasswordPolicy
msDS-PasswordSettingsPrecedence	Yes	Password settings precedence	10
msDS-PSOAppliesTo	No	Multi-valued attribute that holds the DNs of the objects that a PSO applies to	CN=Joe,CN=Users, DC=dc,DC=net
Password Policy-Related Settings			
msDS-PasswordReversibleEncryptionEnabled	Yes	Password reversible encryption status	TRUE
msDS-PasswordHistoryLength	Yes	Password history length	24
msDS-PasswordComplexityEnabled	Yes	Password complexity status	TRUE
msDS-MinimumPasswordLength	Yes	Minimum password length	6
msDS-MinimumPasswordAge	Yes	Minimum password age in days	5
msDS-MaximumPasswordAge	Yes	Maximum password age in days	30
Account Lockout Policy-Related Settings			
msDS-LockoutThreshold	Yes	Lockout threshold	0
msDS-LockoutObservationWindow	Yes	Observation window for lockout of user accounts in minutes	30
msDS-LockoutDuration	Yes	Lockout duration for locked out user accounts in minutes	30

Table 1: *Important AD PSO Attributes*

that determines what AD user accounts or groups the PSO will be linked to. Even though password and account lockout policies can be linked to any AD user, group or computer object, or OU, PSOs are effective only for the user accounts and global groups they are linked to. In addition, PSOs are effective only if your AD domain is in the native Server 2008 domain functional level—which means that all the DCs in your domain must be running Server 2008.

The msDS-PasswordSettings Precedence PSO attribute holds an integer value that is used to resolve conflicts if multiple PSOs are applied to a user or group object. A low value for the msDS-PasswordSettingsPrecedence attribute indicates that the PSO has a higher priority than other PSOs. For example, imagine that a user object has two PSOs linked to it: one PSO that has an msDS-PasswordSettingsPrecedence value of 10 and another PSO that has a value of 40. In this case, the PSO that has the msDS-PasswordSettingsPrecedence value of 10 (the lower value) has a higher rank and will be applied to the user object. If multiple PSOs are linked to a user or group, the logic that Server 2008 uses to determine the resultant PSO is as follows:

- A PSO that is linked directly to the user object is the resultant PSO. If more than one PSO is linked directly to the user object, the PSO with the lowest msDS-PasswordSettingsPrecedence value is the resultant PSO.
- If no PSO is linked to the user object, but PSOs are linked to global groups the user is a member of, Server 2008 compares the msDS-PasswordSettingsPrecedence values of these different global group PSOs. Again, the PSO with the lowest msDS-PasswordSettingsPrecedence value is the resultant PSO.
- If no PSO is obtained from these conditions, the "classic" Default Domain Policy is applied.

To let administrators easily determine the PSO that's ultimately applied to a user, Microsoft added a new attribute called msDS-ResultantPSO to each AD user object. This attribute holds the distinguished name (DN) of the PSO that's applied to a given user.

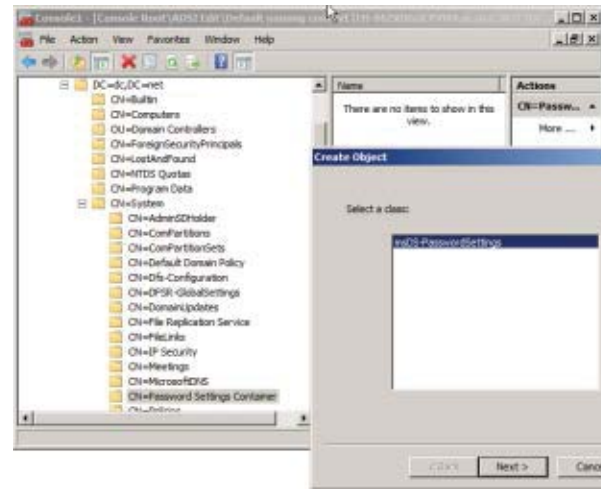

Figure 1: *Using ADSI Edit to create a PSO*

PSO Creation and Configuration Tools

Microsoft doesn't plan to provide a GUI tool or Microsoft Management Console (MMC) snap-in extension to configure fine-grained password policies in the first Server 2008 release. However, you can use existing LDAP query tools such as LDP or LDIFDE, or the MMC ADSI Edit snap-in, to define and configure PSOs. These tools are available on any Server 2008 AD installation. Although these three tools are rather complex, experienced AD administrators should have no problem using them to set the new password policies.

Novice AD administrators, or experienced administrators who simply want to make their jobs easier, might consider Joe Richards' command-line tool called psomgr.exe, or Special Operations Software's Specops Password Policy tool. Specops Password Policy lets you use a special MMC snap-in to configure PSOs from the Windows GUI. Both tools hide the AD complexity behind fine-grained password policies and significantly ease their configuration. You can download the PSOMgr tool from www.joeware.net/freetools/tools/psomgr. The full-featured commercial version of Specops Password Policy is available at www.specopssoft.com/products/specopspasswordpolicy; a free version with limited functionality, called Specops Password Policy Basic, is available at www.specopssoft.com/wiki/index.php/specopspasswordpolicybasic. The full-featured version extends the standard Windows password policy capabilities by adding features such as the ability to disallow the use of user names or certain words in passwords, and automatic user notification of password expiry via email message.

To use ADSI Edit to define a new PSO, start ADSI Edit and connect to the domain where you want to define a fine-grained password policy. Then, navigate to the System\Password Policy Settings container. Right-click the container and select New, Object. In the Create Object dialog box, which Figure 1 shows, select the msDS-PasswordSettings object class, and enter your preferred password and account lockout policy values for the different PSO attributes.

To use LDP to define a new PSO, you must initiate several LDAP commands from the LDP interface. (For information about using LDP, see the Microsoft article "Using Ldp.exe to Find Data in the Active Directory," at support.microsoft.com/kb/224543.) To use the LDIFDE command line to define a new PSO, you must first create an LDF configuration file that specifies the different PSO

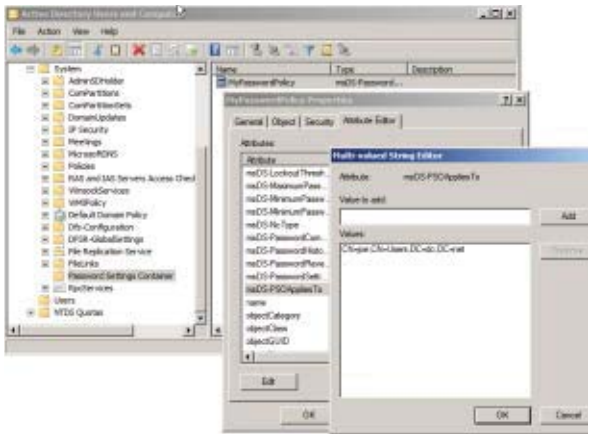

Figure 2: *Modifying the user objects a PSO is linked to from the MMC AQctive Directory Users and Domputers snap-in*

attributes. (For information about using LDIFDE, see the Microsoft article "Using LDIFDE to import and export directory objects to Active Directory," at support.microsoft.com/kb/237677. For more detailed instructions, see the Microsoft article "Step-by-Step Guide for Fine-Grained Password and Account Lockout Policy Configuration," at technet2.microsoft.com/windowsserver2008/en/library/2199dcf7-68fd-4315-87cc-ade35f8978ea1033.mspx?mfr=true.)

When you use the ADSI Edit version that's bundled with Server 2008 to define PSOs, you must enter the four time-related PSO attributes (msDS-MaximumPasswordAge, msDS-MinimumPasswordAge, msDS-LockoutObservationWindow, and msDS-LockoutDuration) in the days:hours:minutes:seconds format. For example, to set a maximum password age of 40 days, you'd enter the value 40:00:00:00.

When you use the ldifde command or an older (pre-Server 2008) version of ADSI Edit to create PSOs, you must enter the values of these attributes in I8 format (i.e., integer represented in 8 bytes). In the I8 format, time must be stored in intervals of -100 nanoseconds. This means that to use LDIFDE or an older ADSI Edit version to set PSO attributes to their appropriate values, you must convert the time you want to set in values in minutes, hours, or days to time values in intervals of 100 nanoseconds, then precede the resultant values with a minus sign (-).

Because the I8 format is difficult to use, I recommend that you use the Server 2008 version of the ADSI Edit tool (or the PSOMgr or Specops Password Policy tools) for defining PSOs. The Microsoft article "Step-by-Step Guide for Fine-Grained Password and Account Lockout Policy Configuration" (technet2.microsoft.com/windowsserver2008/en/library/2199dcf7-68fd-4315-87cc-ade35f8978ea1033.mspx?mfr=true) explains I8 conversion in more detail.

In addition to using ADSI Edit, LDP, LDIFDE, PSOMgr, or Specops Password Policy to link PSOs to users or global groups, you can also use the MMC Active Directory Users and Computers snap-in. To link a PSO to a user or group from this snap-in, open the snap-in and ensure that the Advanced Features view is enabled. (To enable this view, use the Advanced Features option in the View menu.) Then, open the Passwords Settings Container in the System container, right-click the PSO you want to link, and select Properties. In the Properties dialog box, select the Attribute Editor tab, select the msDS-PSOAppliesTo attribute, and click Edit. Finally, in the Edit dialog box, which Figure 2 shows, enter the DN of the user or group you want to link the PSO to. If you don't know the correct DN of a user or group, you can obtain it from the Active Directory Users and Computers snap-in. In the snap-in's details pane, right-click the user or the global security group, select Properties, select the Attribute Editor tab, and view the value of the user's or group's distinguishedName attribute in the Attributes list.

A Valuable Addition

Server 2008's fine-grained password and account lockout policies are a valuable addition to the Windows security management portfolio. Although defining and configuring these policies isn't straightforward in the first Server 2008 release (I strongly advise you to use PSOMgr or the Specops Password Policy tool), the policies do provide a significant level of additional flexibility. For example, Server 2008's fine-grained password policies eliminate the need for organizations to define additional Windows domains or develop special password filters.

InstantDoc ID 97567

Jan De Clercq *(jan.declercq@hp.com) is a member of HP's Security Office and focuses on identity management and security in Microsoft products. He is coauthor of* Microsoft Windows Security Fundamentals *(Digital Press).*

Leveraging Server 2008's Password Policies

1. Windows Server 2003 and Windows 2000 Server password policies let administrators define only one password policy that applies to all user accounts in a domain.
2. Windows Server 2008 introduces fine-grained password policies that let administrators define different password policies for different domain account categories in a single domain.
3. Create Password Settings objects (PSOs) to define new fine-grained password policies.
4. To define and configure the PSOs you created, use an LDAP query tool (e.g., LDP, LDIFDE, ADSI Edit), the PSOMgr or Specops Password Policy tools, or the MMC Active Directory Users and Computers snap-in.

slapd returns only attributes defined in the schema known to OpenLDAP and nothing more:

```
dn: CN=dpuryear,CN=Users,DC=testcorp,D
   C=com
cn: dpuryear
```

Notice that sAMAccountName is not shown, even though it exists in AD and was requested in our ldapsearch command. To access all of the data in AD, you need to install the most recent version of OpenLDAP, OpenLDAP 2.3, which can transparently pass unknown schema, albeit with some minor syntax rules applied so that it can perform filtering.

Using OpenLDAP 2.3 to Pass Unknown Schema

You can install OpenLDAP 2.3 either by compiling the source, or, far easier, by installing it from RPM Package Manager (RPM). After installation, the only configuration change required is to modify pidfile and argsfile because the newer OpenLDAP RPM assumes a different location for those files. Listing 4 shows the code for this.

Now, restart slapd and try ldapsearch again—first using the "cn" in your filter

```
# ldapsearch -x -h localhost -LLL -b
   dc=testcorp,dc=com \
-D cn=dpuryear,cn=users,dc=testcorp,
   dc=com -W \
'(cn=dpuryear)' cn sAMAccountName
Enter LDAP Password:
```

and then the AD-only sAMAccountName:

```
# ldapsearch -x -h localhost -LLL -b
   dc=testcorp,dc=com \
-D cn=dpuryear,cn=users,dc=testcorp,
   dc=com -W \
'(sAMAccountName=dpuryear)' cn
   sAMAccountName
Enter LDAP Password:
```

Success! As you can see from the following output, ldapsearch queried slapd, which in turn queried AD for us:

```
dn: cn=dpuryear,cn=Users,dc=testcorp,
   dc=com
cn: dpuryear

SAMACCOUNTNAME: dpuryear
```

The key difference here is that we now have access to the complete AD schema, including sAMAccountName.

Seamless Access to AD

You should now be able to attach AD to any part of your OpenLDAP directory. You can authenticate your AD users in LDAP applications that use OpenLDAP or even provide access to multiple ADs in your network if they aren't all part of a larger forest already.

InstantDoc ID 98449

Dustin Puryear

(dustin@puryear-it.com) provides expertise in identity management, directory services, and Linux interoperability. He is the author of Best Practices for Managing Linux and UNIX Servers (*Windows IT Pro* eBooks).

When you are looking for just **the right piece**,

look to our new IT and developer resource.

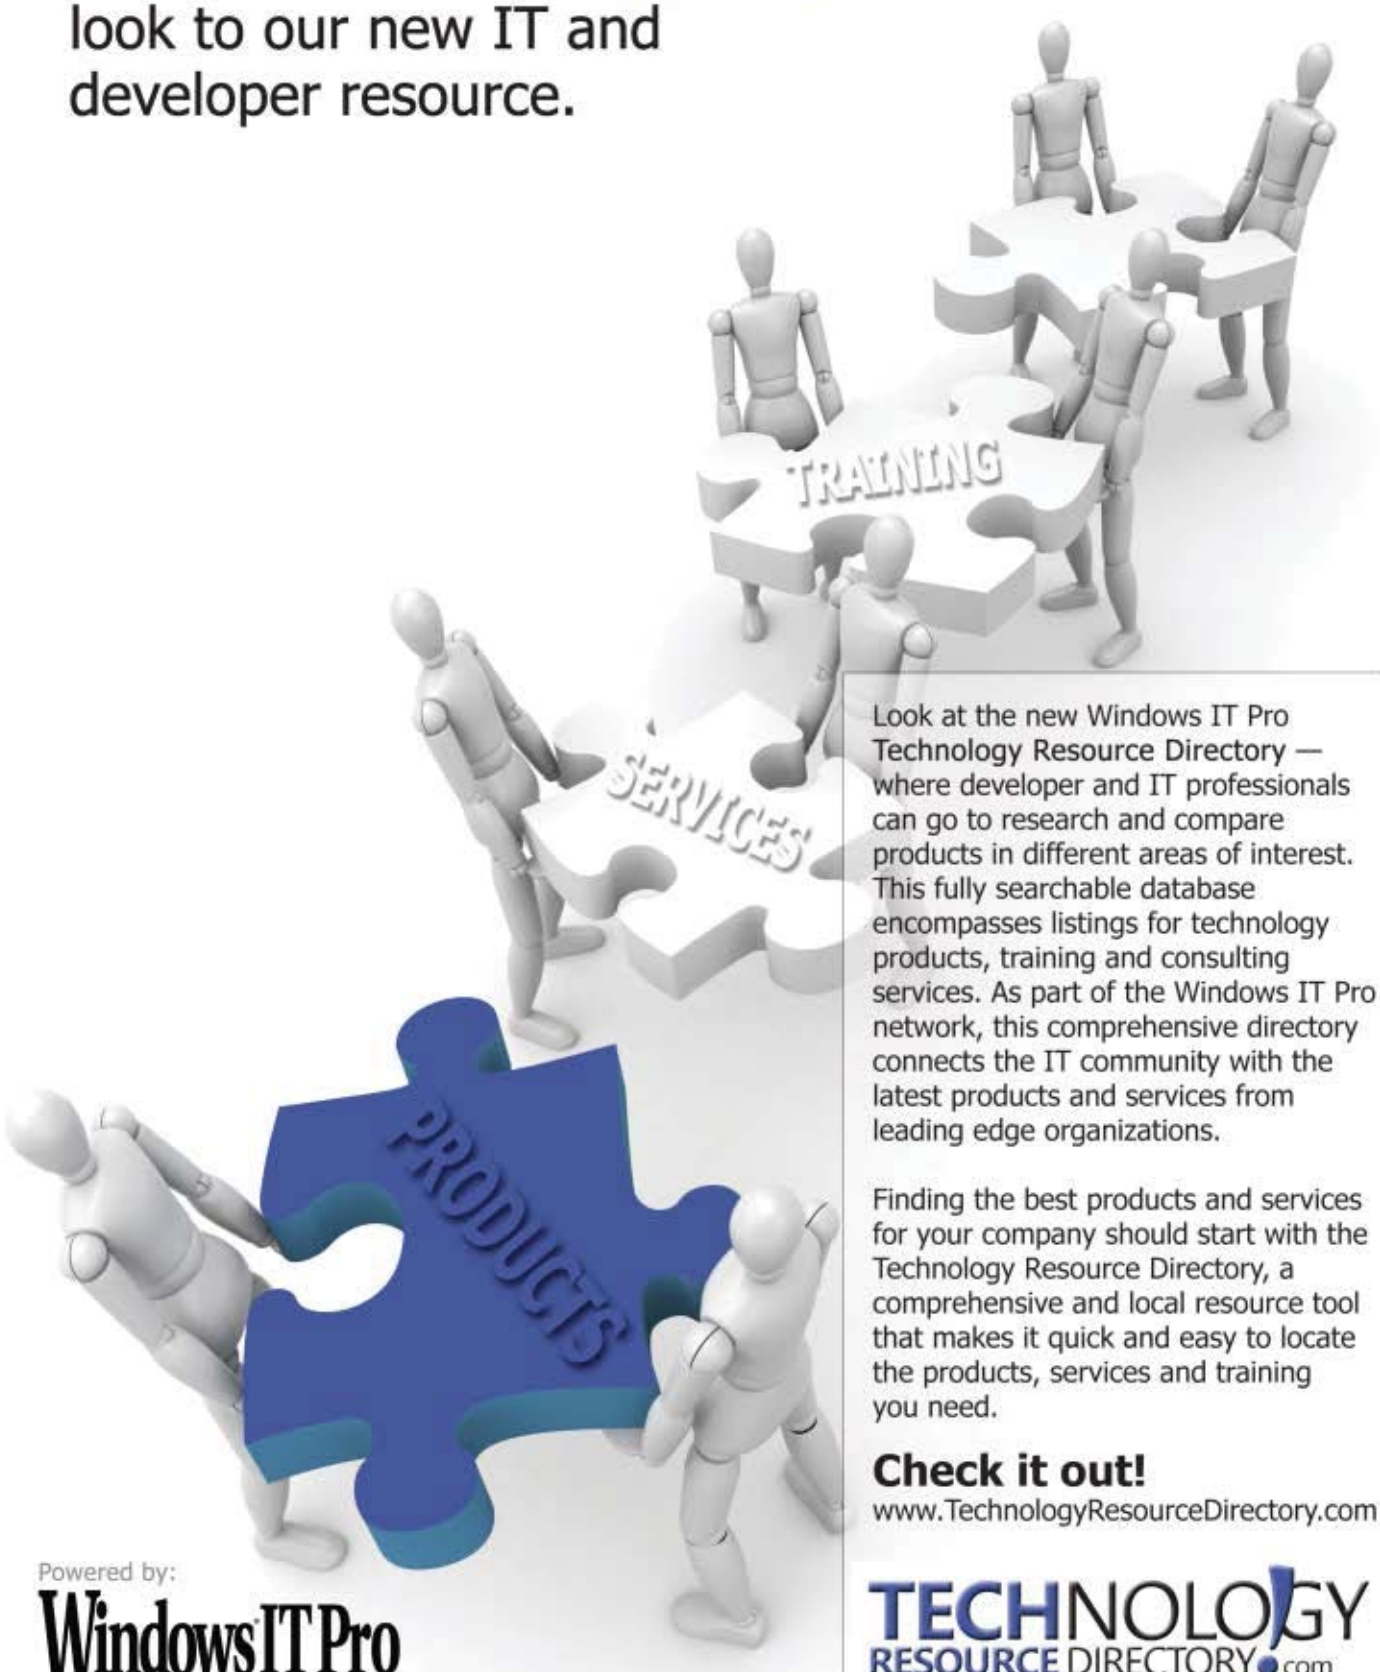

Look at the new Windows IT Pro Technology Resource Directory — where developer and IT professionals can go to research and compare products in different areas of interest. This fully searchable database encompasses listings for technology products, training and consulting services. As part of the Windows IT Pro network, this comprehensive directory connects the IT community with the latest products and services from leading edge organizations.

Finding the best products and services for your company should start with the Technology Resource Directory, a comprehensive and local resource tool that makes it quick and easy to locate the products, services and training you need.

Check it out!
www.TechnologyResourceDirectory.com

Powered by:
Windows IT Pro

TECHNOLOGY RESOURCE DIRECTORY.com

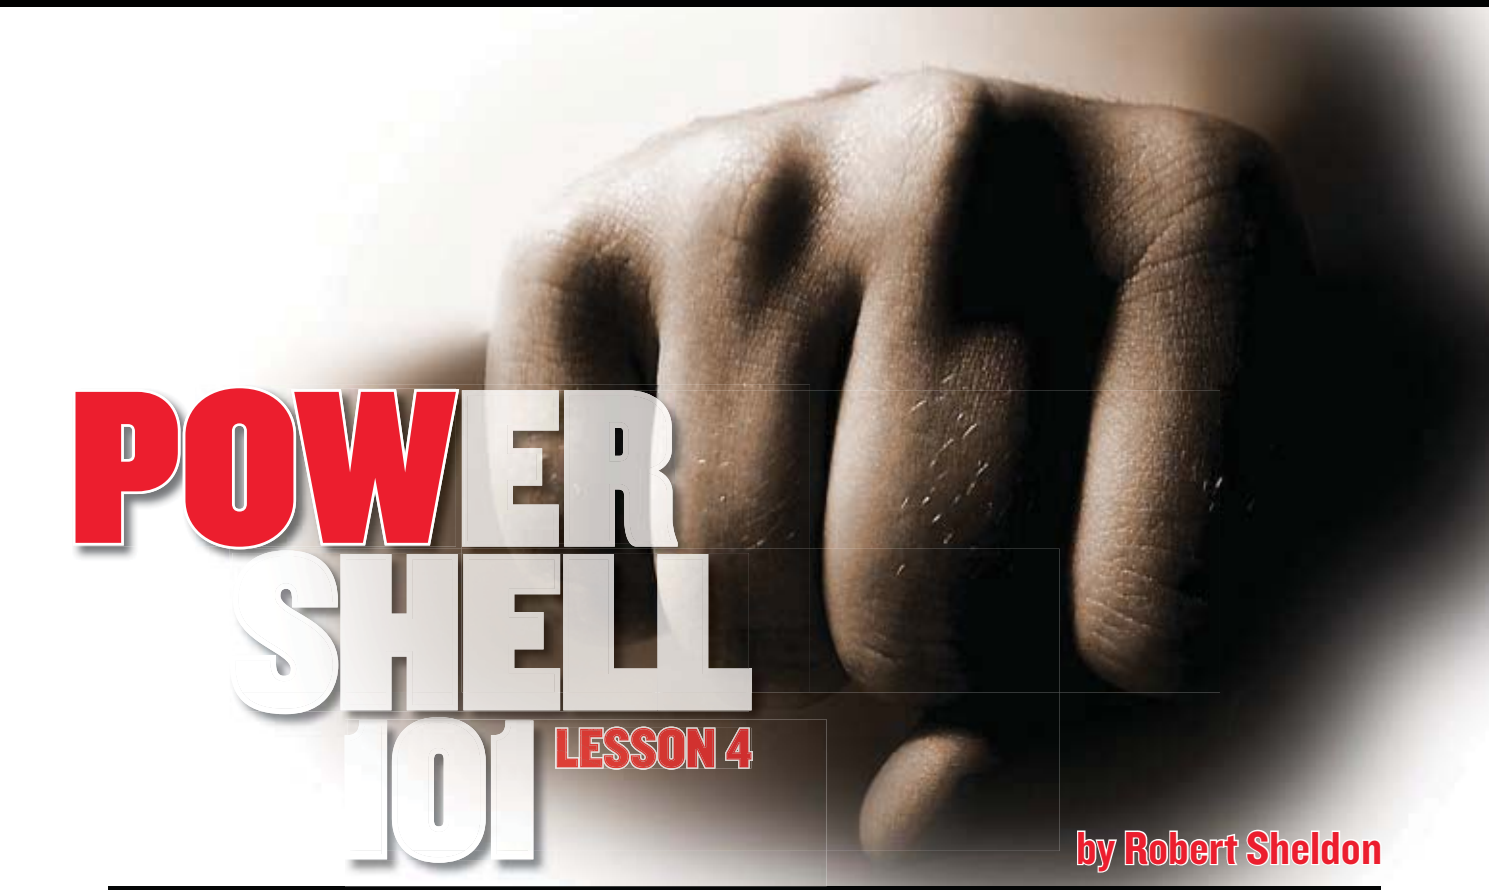

POWERSHELL 101

LESSON 4

by Robert Sheldon

HOW TO PROPERLY USE QUOTES WHEN WORKING WITH STRINGS

Most PowerShell statements include string values. Usually, these strings are passed to cmdlets as arguments. In some cases, the strings are enclosed in single quotes. In other cases, they're enclosed in double quotes. And sometimes they're not enclosed in quotes at all. It's important to understand how to properly handle strings. The rules that govern how to do so are often referred to as *quoting rules*. In this lesson, you'll learn about these rules. Specifically, you'll learn when to enclose string values in quotes and whether to use single or double quotes. In addition, you'll learn how to flag, or *escape*, special characters.

Working with String Values

Whenever you enclose text in quotes, PowerShell treats that text as a string value. So, as long as the text doesn't contain any special characters (you'll learn more about these characters shortly) or reference variables (I'll discuss how to reference variables in strings in the next lesson), you can enclose the text in either single or double quotes. For example, the following statements achieve the same results:

```
Write-Output "String in quotes."
Write-Output 'String in quotes.'
```

In these examples, the Write-Output cmdlet sends a string object down the pipeline or, in this case, directly to the PowerShell console. As you can see in Figure 1, the outputted value is the same for both statements.

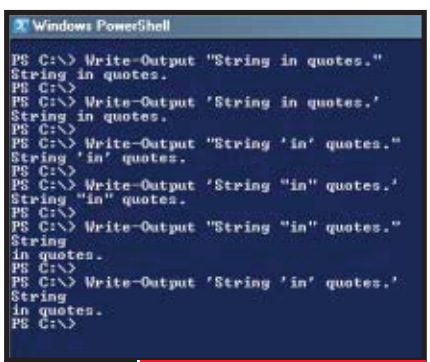

Figure 1: Enclosing strings in single and double quotes

In addition to the Write-Output cmdlet, PowerShell's Out-Host and Write-Host cmdlets output information to the console. Their differences lie in the details. For example, the Write-Output cmdlet sends output down the pipeline to the next cmdlet. When Write-Output is the last cmdlet in the pipeline, the output is displayed in the console. The Out-Host cmdlet sends output directly to the console and offers an optional parameter that lets you view the output one screen at a time, which can be helpful if there is a lot of output. This is the default output cmdlet, so if you don't specify an output cmdlet, Out-Host cmdlet is used. The Write-Host cmdlet also sends output directly to the console. However, Write-Host has two optional parameters that let you change the color of the text or text background, thereby creating a customized console.

For basic quoted string values that you want to output directly to the console window, all three cmdlets behave in similar ways. For example,

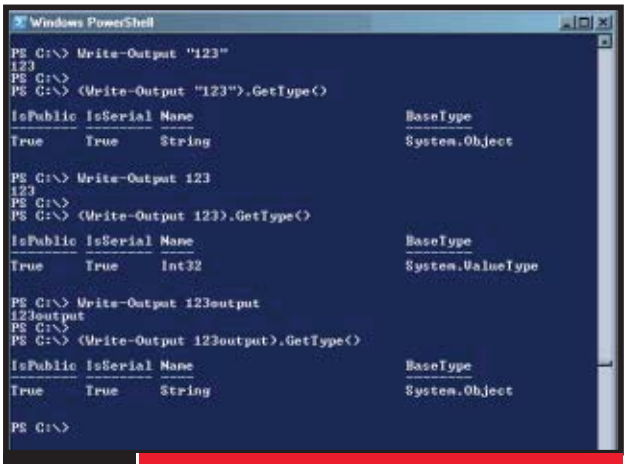

Figure 2: Working with numerical values in strings

the following four commands all display the string in the console window in the same way as in Figure 1:

```
"String in quotes."
Write-Output "String in quotes."
Write-Host "String in quotes."
Out-Host `
  -InputObject "String in quotes."
```

Notice that no cmdlet is specified in the first command. As a result, the Out-Host cmdlet is used.

For many of the examples in this lesson, I use the Write-Output cmdlet because it outputs an object in a way similar to many other cmdlets. This lets me demonstrate different principles about quoted values. Keep in mind, however, that the Write-Output, Out-Host, and Write-Host cmdlets can behave differently in different circumstances. For more information about these cmdlets, see their Help files.

If you want to include quotes within a string, you can use single quotes within double quotes or double quotes within single quotes:

```
Write-Output "String 'in' quotes."
Write-Output 'String "in" quotes.'
```

If you refer again to Figure 1, you'll see that inside quotes in both cases are carried to the output. This isn't the case when you use the same type of quotes throughout the string:

```
Write-Output "String "in" quotes."
Write-Output 'String 'in' quotes.'
```

As Figure 1 shows, the results are quite different. In both cases, the quotes are not displayed and a new line is added. This is because PowerShell interprets the one string as multiple strings and consequently adds a line break. For example, PowerShell interprets *String* as the first string (so it adds a line break after that string), then interprets the rest as a different string. You can use double quotes within double quotes, but you must escape the inside quotes, which I'll describe how to do later.

Whenever you work with quotes, be careful not to mix up the type of quotes or forget to include one. Otherwise, you might get stuck in a loop that continues to prompt you for an entry—but nothing you enter gets you out of the loop. If you run into that situation, press Ctrl+C to return to the command prompt.

Another issue to take into account when defining cmdlet arguments is how PowerShell treats numerical values. As I said earlier, PowerShell treats any values within quotes as strings, even if the value consists of all numbers:

```
Write-Output "123"
```

When you execute this statement (shown in Figure 2), the value returned is a string object, as Figure 2 shows. You can verify the value's type by running the statement:

```
(Write-Output "123").GetType()
```

This statement uses the GetType method of the object's type, which in this case is String, or more specifically System.String, also shown in Figure 2. For more information about System.String and GetType, see the sidebar "Getting and Using the System.String Object's Members."

If you don't enclose a numerical value in quotes, PowerShell treats the value as a numerical object. For example, the following statement returns an integer object:

```
Write-Output 123
```

Again, you can verify the object's type by using the GetType method:

```
(Write-Output 123).GetType()
```

As you can see in Figure 2, the object's type is Int32.

If a value includes both numbers and letters, PowerShell treats it as a string, whether or not it's in quotes. For example, the following statement returns a string, as verified by the second statement:

```
Write-Output 123output
(Write-Output 123output).GetType()
```

Once again, Figure 2 shows the output of these statements.

In most cases, you can omit quotes if your argument is a string with no embedded spaces. For example, the following three statements use the Set-Location cmdlet to set the working folder to the C drive:

```
Set-Location C:\
Set-Location "C:\"
Set-Location 'C:\'
```

Now suppose you want to change the working folder to C:\Documents and Settings:

```
Set-Location `
  C:\Documents and Settings
```

This statement generates an error because it doesn't know what to do with the tokens (words) after the first space. As you can see

Learning Path

WINDOWS IT PRO RESOURCES

To read the previous "PowerShell 101" lessons, go to:

"PowerShell 101, Lesson 1," InstantDoc ID 97742

"PowerShell 101, Lesson 2," InstantDoc ID 97959

"PowerShell 101, Lesson 3," InstantDoc ID 98177

SCRIPTING PRO VIP RESOURCES*

If you're beyond the basics, check out:

"Enhance PowerShell's Syntax Display," InstantDoc ID 98088

"Working with Shortcuts in Windows PowerShell," InstantDoc ID 97837

* Need to be a Scripting Pro VIP subscriber. Go to www.scriptingprovip.com for information.

Getting and Using the System.String Object's Members

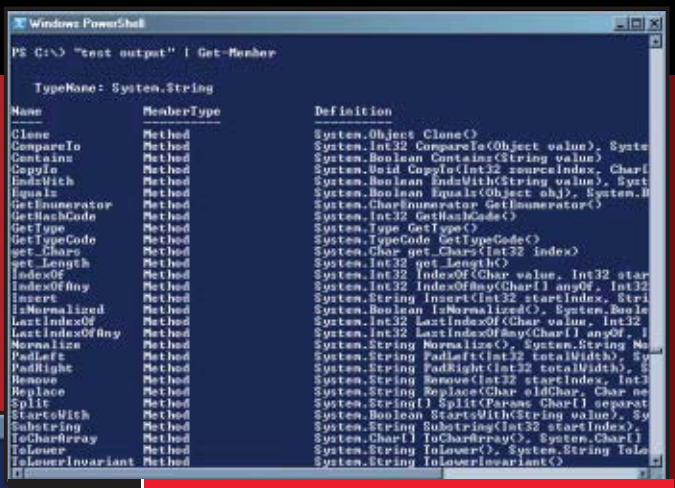

Figure A: Retrieving the members of a System.String object

PowerShell treats all strings as System.String objects. This means that you can use the rich set of methods and properties available to those objects.

As shown Lesson 3, the Get-Member cmdlet retrieves an object's members as the object is passed down the pipeline. Because a string is passed down as an object, you can use Get-Member for that string. For example, the following statement retrieves the object members for the string *test output:*

```
"test output" | Get-Member
```

As you can see in Figure A, a string object supports numerous methods, including Substring and GetType. If you were to scroll down, you would also find the Length property, which provides the number of characters in the string.

Now suppose you want to learn more about the Substring method of the System.String object. You can again use Get-Member to retrieve the information you need. You specify the Substring method as an argument to Get-Member, then pipeline the results to the Format-List cmdlet:

```
"test output" |
Get-Member Substring |
Format-List
```

Figure B shows the statement's result. Notice that the definition includes details about how to use the method. This is the method's syntax. The syntax specifies that you can choose one of two approaches when calling this method. Those approaches are:

```
System.String Substring(Int32 startIndex)
System.String Substring(Int32 startIndex, Int32 length)
```

In the first approach, you provide the target string and an integer that specifies the position in which you want the substring to start. The method will then return a substring that begins at this position and ends at the end of the string. For example, if the target string is *test output* and you want to return a substring that starts at position 5, you'd use the statement

```
("test output").Substring(5)
```

which would return the substring *output.* As this statement shows, you must enclose the target string in parentheses, then add a period followed by the method name and the parameter enclosed in parentheses.

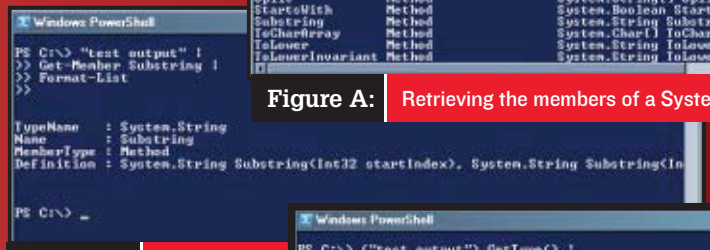

Figure B: Retrieving details about the Substring method

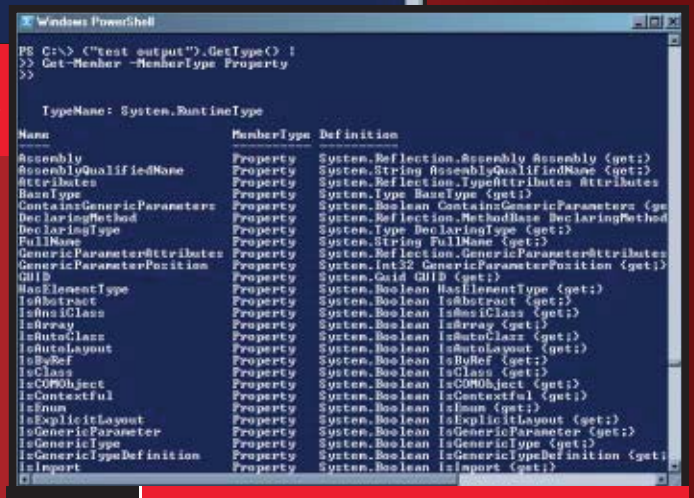

Figure C: Retrieving the GetType method's properties

In the second approach, you provide the target string and the substring's starting position and length. For example, suppose you want to return a substring that starts at the beginning of the string (position 0) and is four characters long. You'd use the statement

```
("test output").Substring(0,4)
```

which would return the substring *test.*

When you call a method such as Substring or GetType, it generates its own object that can be passed down the pipeline. As a result, you can use Get-Member to retrieve a list of those object's members. For example, the GetType method returns a System.RuntimeType object, and that object supports numerous methods and properties. The following statement returns a list of properties available to the System.RuntimeType object:

```
("test output").GetType() |
Get-Member -MemberType Property
```

As the statement shows, you first enclose the string in parentheses, add a period followed by the method name, which in this case is GetType. You then pipe the object returned by GetType to Get-Member. Notice that I've used the -MemberType parameter to display only the returned object's properties. Figure C shows the results.

InstantDoc ID 98448

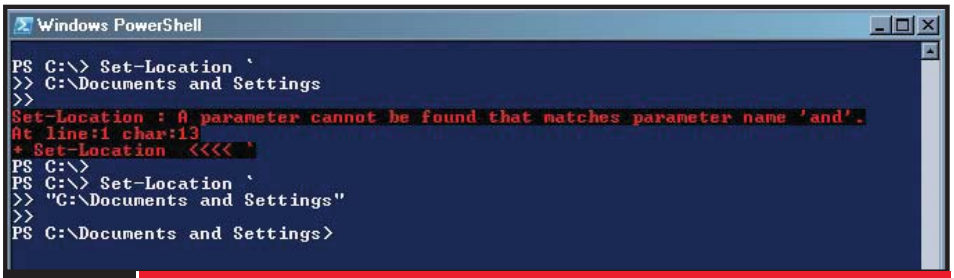

Figure 3: Handling arguments improperly results in an error message

Table 1: PowerShell's Special Characters

Special Character	Description
`0	Inserts a Null value
`a	Sends an alert (bell or beep) to the computer's speaker
`b	Inserts a backspace
`f	Inserts a form feed
`n	Inserts a new line
`r	Inserts a carriage return
`t	Inserts a horizontal tab
`v	Inserts a vertical tab
`'	Inserts a single quote
`"	Inserts a double quote

in Figure 3, the parser interprets "and" as a parameter, and because there's no parameter by this name, the parser generates an error. You can easily fix this problem by enclosing the entire argument in quotes:

```
Set-Location `
  "C:\Documents and Settings"
```

Now when you run this statement, it changes the working folder, as shown in Figure 3.

One other important issue when working with strings is how to reference variables within a quoted string. If you enclose a string in double quotes, the variable's value is used. If you enclose a string in single quotes, the literal value is used. As I mentioned earlier, I'll cover variables and describe how to reference them in strings in the next lesson. However, if you're anxious to learn more about using variables in your strings now, refer to the about_quoting_rules Help file.

Escaping Special Characters in Strings

Up to this point, the arguments you've seen in the examples could have taken single or double quotes. As a result, it would appear that there's no difference between the two. However, there's one very important difference: Single quotes always treat a string literally, whereas double quotes allow you to escape special characters within the text. A special character is one that, when preceded by a backtick (`), takes a specific action that it would not have taken without the backtick. Table 1 lists PowerShell's special characters.

The best way to explain this concept is through a few examples. In the following statement, several characters have been escaped in order to change how the text is displayed:

```
Write-Output ("`n`tText includes" + `
  "`n`t`"escaped`" characters.`n")
```

The first escaped character (the one preceded by the first backtick) is *n*, which in this context inserts a new line. The next escaped character is *t*, which inserts a tab. Note that the backtick at the end of the first line isn't being used as an escape character but rather as a continuation character (see Lesson 2). In the second line, `*n* and `*t* are used several more times. In addition, backticks precede the double quotes surrounding the word *escaped*. As a result, the double quotes appear in the output. If you refer to Figure 4, you'll see how the new lines, tabs, and double quotes appear. For more information about escaping characters, see the about_escape_character Help file.

If you try to escape characters in a string enclosed in single quotes, the backtick and special characters have no effect on the output, other than to be treated literally. For example, the statement

```
Write-Output '`tindented`n`twords'
```

returns the exact string as originally typed, as Figure 4 shows. Note that, in pre-release versions of PowerShell, you could escape characters inside single-quoted strings, and PowerShell would correctly parse them. So, you might come across material that says that this is how PowerShell handles single-quoted strings.

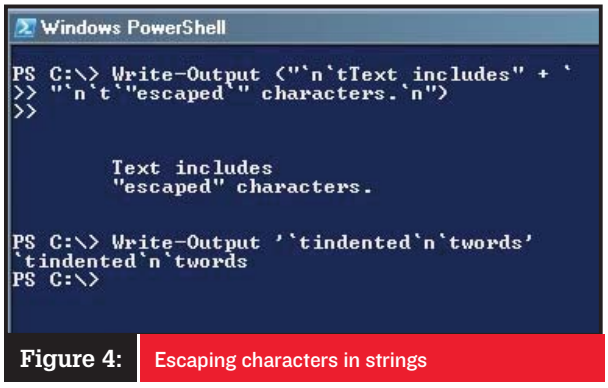

Figure 4: Escaping characters in strings

Moving Forward

In most cases, string values play an important role in creating PowerShell statements. The better you understand how to work with strings, the more effective your statements will be. I encourage you to spend time practicing the various ways to use strings. Try enclosing them in single and double quotes, then try running the commands without the quotes. And don't forget to practice escaping special characters.

InstantDoc ID 98447

Robert Sheldon (contact@rhsheldon.com) is a technical consultant and the author of numerous books, articles, and training material related to Microsoft Windows, various relational database management systems (including SQL Server), and business intelligence design and implementation. He is also the author of the novel *Dancing the River Lightly*.

3 Info-packed eLearning seminars for $99!

EXCHANGE 2007 Mastery Series—Part 2

Hosted by Windows IT Pro

WHEN

May 29, 2008—11:00 AM, EDT

WHERE

On your computer

COST

$99/registrant for 1, 2, or all 3 live online sessions, and includes access to the archived versions

SESSIONS

- **Mailbox High availability in Exchange 2007: Learn the Pros and Cons of Your High-Availability Options**
- **Transport Rules: See Real-World Examples You Can Implement in Your Environment**
- **PowerShell: Get Started with Basic Commands—You Don't Need a PhD in Rocket Science**

RESERVE A SEAT by going to:

www.windowsitpro.com/go/elearning/MasteringExchange2007

SPEAKER

Mark Arnold

MCSE+M, Microsoft MVP

Mark Arnold is a senior technical architect for Anix, a UK-based storage integrator, where he solves storage and compliance problems for his clients by using Microsoft Exchange as a key component in SAN and NAS deployments. He's also a regular contributor to Microsoft's "Industry Insiders" TechNet program and is active on Exchange newsgroups and forums.

ABOUT THE SESSIONS

Mailbox High Availability in Exchange 2007: Learn the Pros and Cons of Your High-Availability Options

Exchange 2007 now has several acronyms for high availability—LCR, CCR, SCR—not including anything you can do with your storage or CDP solutions. Which method is best for you? How can you implement a mix of options to make your environment highly available at a price point that doesn't break the bank?

Transport Rules—Real-World Examples That Can Help Your Environment

What are transport rules and how do they help you administer your Exchange environment? You can find many complicated—and largely useless—examples on the Internet. We'll show some interesting things you can do with message flow and give you real-world examples.

PowerShell—The Things You Need That Don't Involve a PhD in Brain Surgery or Rocket Science

Some Exchange admins resist PowerShell because they think they can complete tasks quicker through the GUI. But we'll present some useful, quick, and readily repeatable PowerShell commands that will make your job easier rather than your hair grayer.

REGISTER TODAY — seats are limited,

to allow lots of live Q&A at the end.

For more information, or to register, go to:

www.windowsitpro.com/go/elearning/masteringexchange2007

Windows IT Pro

Introducing an integrated approach to complete SharePoint protection and management

DocAve™ Software for SharePoint

Changing the way Administrators manage SharePoint

FREE 30 DAY TRIAL
Download at
www.avepoint.com

SharePoint management made simple.

Now you can control and manage the back-end of all your SharePoint environments from one place. DocAve is the only truly integrated, easy-to-use software that offers a complete set of SharePoint backup, recovery, and administration tools. One solution, with many mix-and-match functions, now gives you power like never before.

Complete SharePoint protection.

With item-level backup and full-fidelity restore, DocAve allows for fast recovery of business critical documents and content. Complete SharePoint platform backup allows for quick and painless recovery of the entire system during a disaster. With DocAve, you'll have complete confidence in your SharePoint environment.

AvePoint

Call 1-800-661-6588 or visit www.AvePoint.com for more information or to download a free trial.

Office & SharePoint PRO
officesharepointpro.com

INTEGRATING Exchange Server 2007 and SharePoint Server

Combine Exchange and Office to improve collaboration scenarios

by Brien M. Posey

We all know that Microsoft Office Outlook 2007 acts as the primary interface to Microsoft Exchange Server 2007. You might not realize, though, that other Microsoft Office products are also designed to integrate easily with Exchange Server. The best example is Microsoft Office SharePoint Server 2007 (MOSS 2007), the workplace collaboration and content management platform. When you configure MOSS 2007 to work with Exchange 2007, the users in your organization will be able to easily share documents over a corporate intranet.

Quick Introduction to MOSS 2007

MOSS lets you create an internal Web site (an intranet site) for use by your company's employees. An intranet site can be used to display corporate announcements and provide access to the corporate directory, but you can set up such sites manually without using MOSS. MOSS's true value is in letting you establish a document library on your intranet site that allows users to check out, modify, and return documents. Using MOSS permissions, you can control which users are able to read or modify a document.

MOSS also lets users create additional Web sites very easily. For example, a group of employees working together on a project can create an intranet site dedicated to that project, then use that site to share project-related documents, post a calendar of project-related milestones, and provide contact information for those involved in the project.

As you can imagine, MOSS is a fairly complex product, but it's surprisingly intuitive. After all, it was designed so that even end users can create complex sites.

Learning Path

WINDOWS IT PRO RESOURCES

For more details about integrating SharePoint with Outlook:

"SharePoint Integration with Outlook 2007, part 1," InstantDoc ID 95919

"SharePoint Integration with Outlook 2007, part 2," InstantDoc ID 96154

"SharePoint Integration with Outlook 2007, part 3," InstantDoc ID 96384

"Shared Calendars with SharePoint and Outlook 2007," InstantDoc ID 96663

To learn more about SharePoint Web Parts:

"SharePoint Portal Server Visual Elements," InstantDoc ID 93399

MICROSOFT RESOURCES

Microsoft Office SharePoint Server Home Page, office.microsoft.com/en-us/sharepointserver/FX100492001033.aspx

SharePoint Server 2007 Web Parts Learn More, microsoft.com/downloads/details.aspx?familyid=760989l0-ead4-4772-bfb0-aedd7248647a&displaylang=en&tm

A Few Prerequisites

Before I show you how MOSS interacts with Exchange Server and what this combination can do for your organization, I need to share some assumptions that I make in this article. I assume that you're running Exchange Server 2007 and that you have at least one client access server deployed. I also assume that your client access server is configured to act as a front end to your Exchange organization and that the client access server role is not installed directly on a mailbox server.

Another prerequisite is that you need to install MOSS on a dedicated server within your perimeter network. The MOSS server must be able to communicate with your mailbox servers, but for performance and security reasons, you shouldn't install MOSS on a server running Exchange.

And finally, I'm assuming that you have an established Exchange organization, that you've just installed MOSS, and that you're starting from scratch.

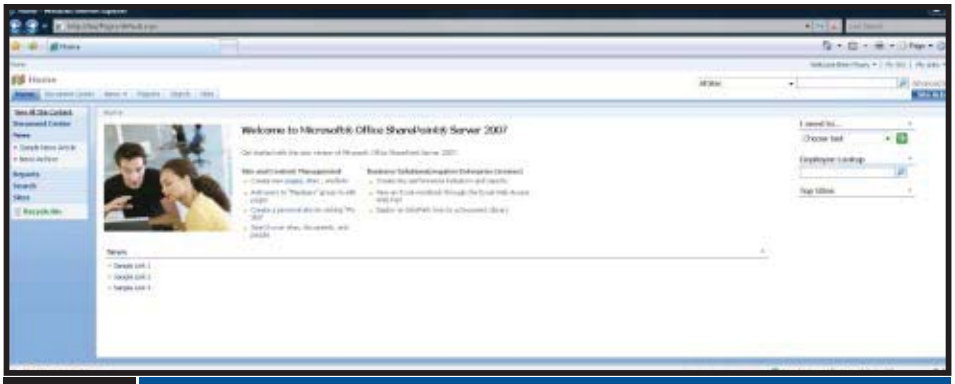

Figure 1: The default SharePoint Web site

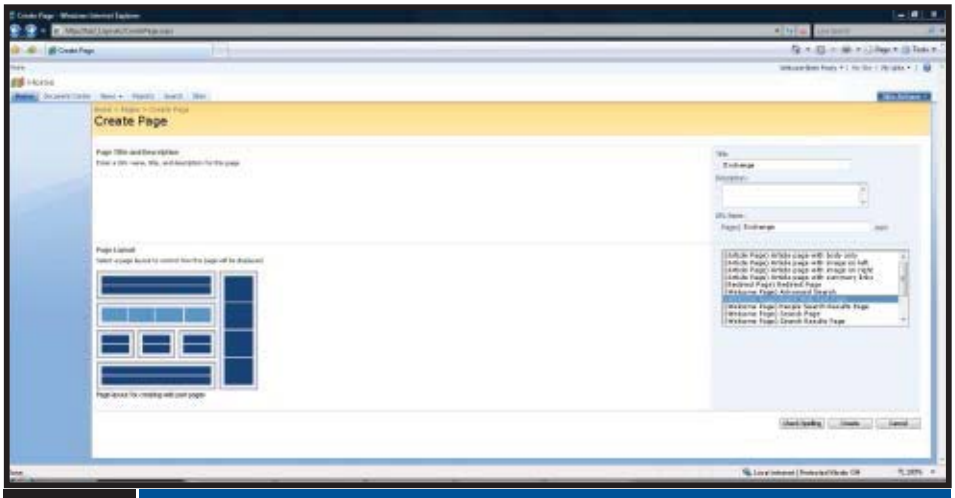

Figure 2: Entering a page name and selecting a template

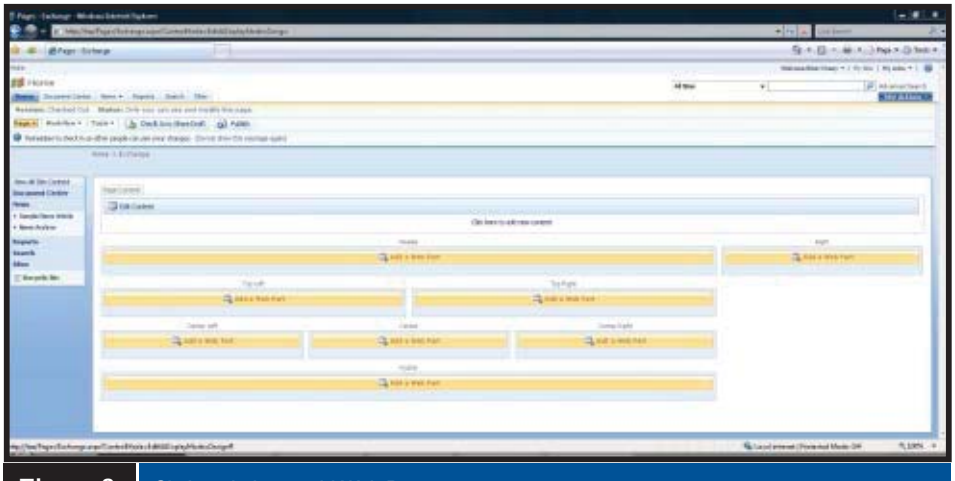

Figure 3: Clicking links to add Web Parts

Creating a SharePoint Web Site

Now that the prerequisites are taken care of, it's time to create a SharePoint Web site that interacts with Exchange Server. First you need to open the default SharePoint site by starting Microsoft Internet Explorer (IE) and entering the URL http://*your_server*/pages/default.aspx, where *your_server* is the NetBIOS name of your SharePoint server. Upon entering this URL, you'll see the default SharePoint Web site displayed in IE, shown in Figure 1. Although the default site has nothing to do with Exchange Server, you can integrate SharePoint Server and Microsoft Outlook Web Access (OWA) into a SharePoint site. This allows you to take advantage of one-stop shopping. You don't have use a separate Web site to access your Exchange mailbox; you can do it directly through the SharePoint site.

Let's add a user's Inbox and Calendar to the default SharePoint site. Under Site and Content Management, click the *Create new pages, sites, and lists* link. You'll see a screen like the one in Figure 2, which lets you work with Web Parts to create a SharePoint Web site. Because the average user doesn't know how to write ASP.NET or HTML code, MOSS includes dozens of predefined Web Parts, which are blocks of code that accomplish a specific task. You plug Web Parts into predefined templates to create Web pages—the entire process can be completed in a matter of minutes. You can also develop your own Web Parts. For more information

on MOSS and Web Parts, see the Learning Path on page 57.

On the Create Page page, enter a title and a description for the page you want to create. From the list box on the right-hand side of the page, select a template for the page layout. Because we'll be adding Web Parts to the page, I've selected the Blank Web Part Page template. Make a note of the URL assigned to the page.

Next, click Create, and you'll see the screen shown in Figure 3. This template contains several links that you can click to add a Web Part to a part of the screen, such as Header, Footer, Left, Right, or Center.

Click the Add a Web Part link beneath the screen section labeled Top Left, and you'll see the list of OWA-related Web Parts. Scroll through this list, select the check box next to the My Inbox Web Part, then click Add. Repeat the process to add the My Calendar Web Part to the Center portion of the Web page. (You can add other Web Parts if you wish.) The template screen should now look like Figure 4.

Notice that each Web Part in Figure 4 contains a link that you must click to configure the Web Part. This link is the Edit link. It doesn't appear until a Web Part has been added. When you click these links, the only information you need to provide is the name of your Exchange server. Enter the name as the URL to your OWA server in the top portion of the My Inbox section (which isn't visible in Figure 5, because the contents of My Inbox have been scrolled down). You can use the various fields on the page shown in Figure 5 to customize the size and appearance of the Web Part.

After you've configured the Web Parts and clicked OK, SharePoint displays the OWA sign-in screen in place of each OWA Web Part, as Figure 6, page 58, shows. Keep in mind, however, that you're viewing the template, not the page itself. To view the actual page you've created, enter the page's URL. For example, I named my sample page Exchange, so the URL would be HTTP://*server_name*/pages/Exchange.aspx. When you connect to the page you've created, you'll see the OWA logon prompts. After you log on, you'll see a page like the one in Figure 7, page 58, where you can see that the unused placeholders from the template aren't displayed. Only the Web Parts that you've added and configured are shown.

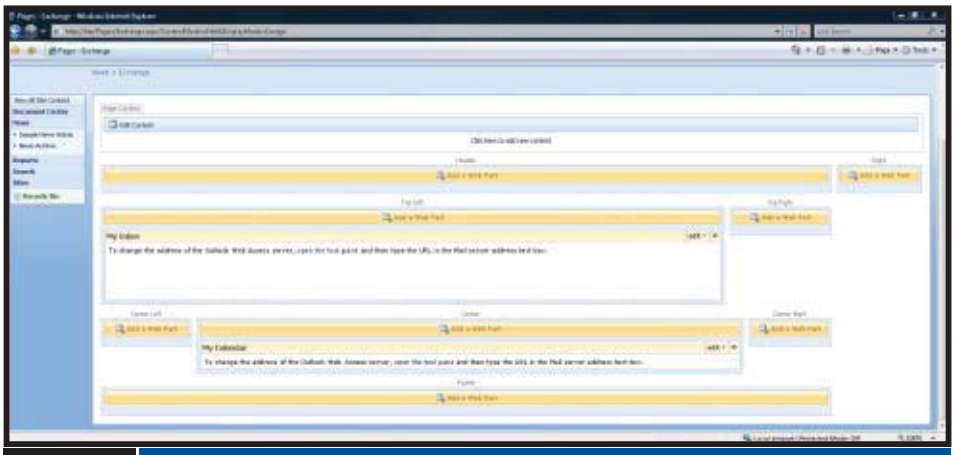

Figure 4: Template with added Web Part

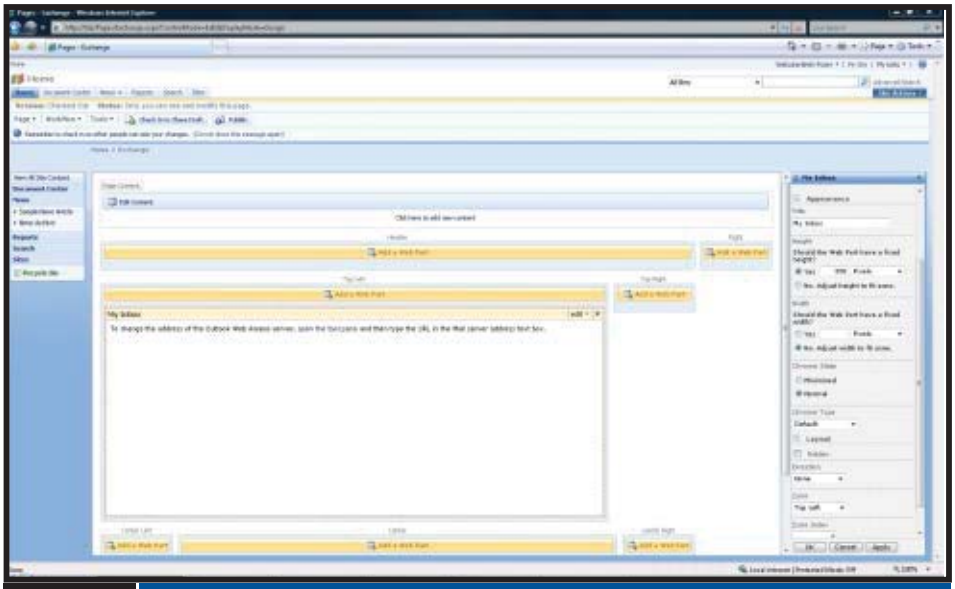

Figure 5: Customizing the Web Part's appearance

The SharePoint Document Library

I mentioned earlier that a primary feature of SharePoint is its document library, which acts as a repository for all document files users create. One interesting thing about the document library is that it's accessible through OWA.

Microsoft introduced this functionality to solve a common problem: Instead of attaching a document to an email message, users often provide a link to the document in the message. In previous versions of Exchange Server, the link worked fine as long as the message recipient was logged on to the domain and was using Outlook to view the message. If the recipient was working outside the office, though, and was using OWA to view the message, the link to the document didn't work.

Microsoft has corrected this problem in Exchange 2007. Now, when a user clicks a link to a document through OWA, the Exchange server sends a request on the user's behalf to the SharePoint server that has the document. After the document is retrieved, it's sent to the user. Depending on how Exchange has been configured and on the document type, users can open the document in a Web browser

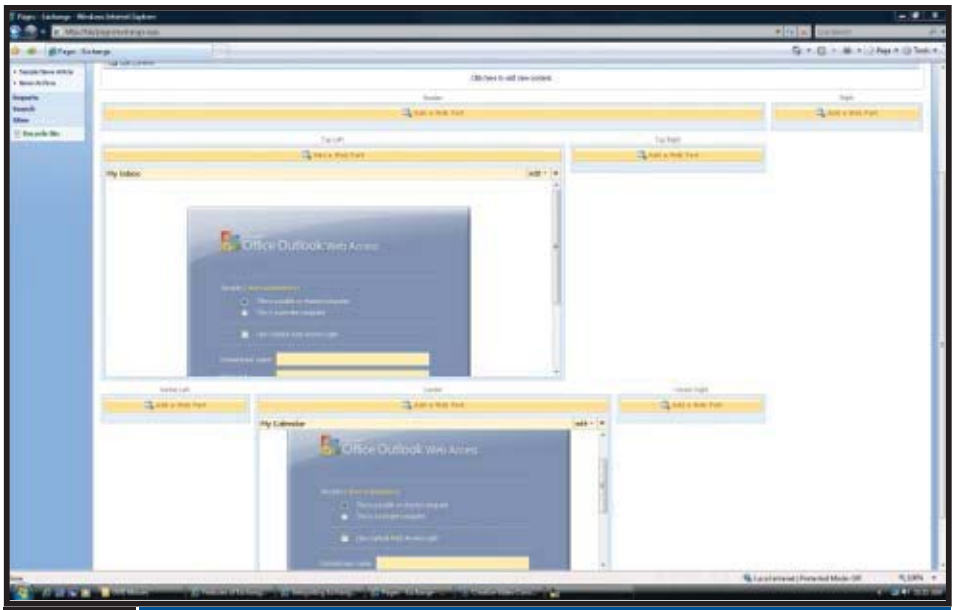

Figure 6: Template with OWA sign-in screens

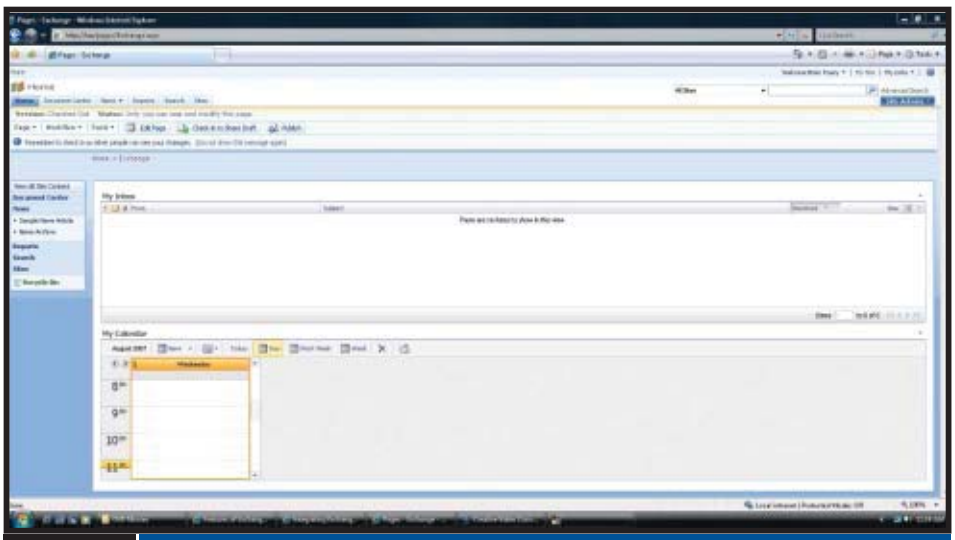

Figure 7: The finished page

or through the application that's associated with it. Incidentally, this process also works if the document is located on a traditional file server.

Begin the process of making the document library accessible through OWA by opening the Exchange Management Console and navigating to Server Configuration, Client Access. Select the client access server you want to configure from the details pane, then right-click OWA (Default Web Site) in the work pane and select Properties from the context menu.

On the OWA (Default Web Site) Properties sheet, click the Remote File Servers tab. Click Configure and enter the domain suffixes you want to treat as internal (so that the domain is trusted; SharePoint allows servers within trusted domains to be accessible). For example, my public domain name is brienposey.com, but the domain name used internally by my production network is production.com, so I'd enter production.com. Verify that the Unknown Servers option is set to Block, which prevents users from accessing unauthorized servers through OWA. Finally, click Allow, and enter the Fully Qualified Domain Name for each SharePoint or file server you want your users to be able to access through OWA.

When users log on to OWA, a prompt asks whether they are using a private or a public computer. Exchange 2007 lets you configure remote file access differently depending on how the user responds to that prompt. Keep in mind, though, that users are on the honor system (scary thought, isn't it?); there's no way to verify whether the user is using a public or private computer.

The OWA (Default Web Site) Properties sheet includes a Public Computer File Access tab and a Private Computer File Access tab. The options on both tabs are identical, letting you configure file access differently depending on which type of computer the user claims to be using. On both tabs, you select a check box to enable direct file access. You can enable file access for Windows file shares, SharePoint, or both. You can also enable WebReady Document Viewing, which lets users view documents in a Web browser even if the application in which the document was created isn't installed on their computer.

To use WebReady Document Viewing, Exchange must have a document converter for the specific file type. Office 2007 document converters are included with Exchange Server 2007 SP1.

Create Custom Collaboration Solutions using Exchange and MOSS

Now that you've seen how Exchange 2007 interacts with MOSS, by enabling OWA Web Parts through SharePoint, and by allowing access to documents stored in a SharePoint document library through OWA, you're ready to start planning custom collaboration solutions for your organization.

InstantDoc ID 98135

Brien M. Posey

(www.brienposey.com) is the vice president of research for Relevant Technologies. He writes technical content for a variety of publications and Web sites.

Office & SharePoint PRO
officesharepointpro.com

CONTROLLING SharePoint Access

A little knowledge about authorization and authentication processes will help you control which users have access to which SharePoint resources

by Kevin Laahs

SharePoint 2007 provides many mechanisms for controlling user access to resources. It provides flexibility in authenticating and identifying users as well as granularity in authorizing what users can do once they've been identified. Understanding SharePoint's end-to-end security model and the major components of the authentication and authorization architecture will help you design a security model to suit your business needs. In this article, I use the term SharePoint to refer to both Windows SharePoint Services (WSS) 3.0 and Microsoft Office SharePoint Server (MOSS) 2007, and I'll call out specific product names when necessary.

The Site Framework—Web Applications and Site Collections

Understanding the authentication and authorization process requires an appreciation of what I term the site framework—in particular, Web applications and site collections. SharePoint is a Web-based application, and it's Microsoft IIS running on Web front-end servers that initially processes user requests. Ultimately these requests come in the form of a URL (e.g., http://friends.laphroaig.com, http://islay.com/sites/ardbeg, or http://islay.com/sites/friends/laahs/myplot.doc), and therefore SharePoint-specific processing must occur based on the incoming URL.

IIS can support multiple Web sites, and each site is uniquely defined using a combination of IP address, port number, and host header. When you want SharePoint to handle incoming URLs, essentially you extend an IIS Web site turning it into a SharePoint Web application, whereupon SharePoint handles the URL requests. (In WSS 2.0 and SharePoint Portal Server—SPS—2003 terminology, a Web application is known as a virtual server). And just as IIS can host multiple Web sites, so too can you have multiple Web applications.

Do you need multiple Web applications? That's a question you need to ask as part of your deployment planning, and many factors will influence this decision, such as differing security needs, need for content isolation, and namespace requirements. From a security standpoint, it's the Web applications that essentially control the authentication method that SharePoint uses for a particular URL namespace. Given that you can have different Web applications pointing to the same content, you can use different authentication schemes for different scenarios (e.g., accessing a SharePoint site from the Internet as opposed to accessing it from the Intranet).

In a SharePoint 2007 farm, all Web applications exist on all Web front-end servers with their definition being automatically replicated by the topology platform service so that the

same authentication method is used consistently, regardless of which Web front-end server receives your initial request. (In WSS 2.0 and SPS 2003, you had to manually clone any virtual servers across physical Web front-end servers).

Web applications host one or more site collections. As the name suggests, this is a collection of sites consisting of one top-level site and one or more subsites. All the sites in a site collection will be part of the same URL namespace from the top-level site down. So, for example, if http://laphroaig.com/sites/friends is your top-level site, then a subsite called plots will have the URL http://laphroaig.com/sites/friends/plots. Each site within the collection then houses its own content in terms of items contained within lists and libraries, both of which can contain any number of folders. Do note that various resources, such as site-wide columns and quotas, are shared among the different sites in a collection. Also, site collections are essentially the unit of administration, are wholly contained within a single content database, and can be backed-up and restored into a different content database, even within a different SharePoint farm. You can see in Figure 1 how site collections are organized within a farm, and Figure 2 provides a sample deployment that supports multiple namespaces through multiple Web applications.

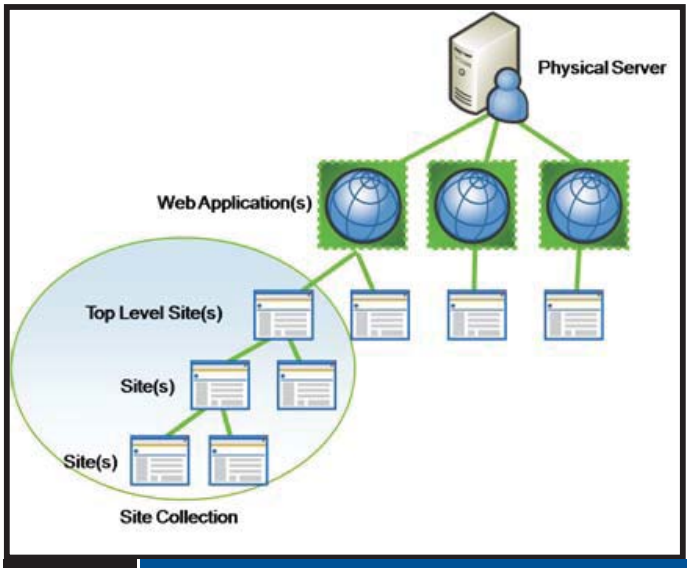

Figure 1: Organization of SharePoint site collections within a farm

Granting users the rights to perform different actions on the resources within a site collection is what we mean by authorization. However, before you can do this you need to be able to identify who someone is via authentication.

Authentication

IIS manages the authentication process for SharePoint 2007 and, unlike WSS 2.0/SPS 2003, SharePoint 2007 no longer limits you to standard Windows authentication methods. Most current SharePoint implementations use Active Directory (AD) as the identity management system against which users are authenticated, but with the support of ASP.NET pluggable authentication, SharePoint 2007 has now opened the door to any number of authentication methods and, by implication, to any number of identity management systems. For example, you can now choose to verify who someone is against a simple list of usernames and passwords held in a file or database or against an enterprise directory system of your choice.

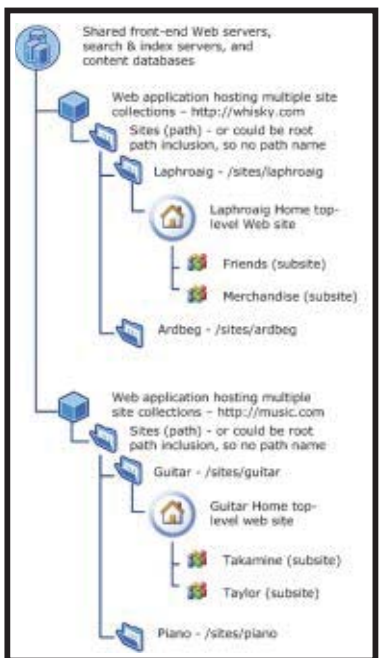

Figure 2: Sample deployment that supports multiple namespaces through multiple Web Applications

This functionality requires something called a membership provider, which can query the specific repository and validate whether a particular user exists based on a supplied username and password. The username and password are collected from the requesting user via a form and passed to the membership provider. ASP.NET 2.0 (upon which SharePoint is built) ships a SQL membership provider, and MOSS provides an LDAP one, but you can also roll your own if needs dictate. Optionally, you can implement a role provider that can define groups of users within your overall membership, and these groups, as well as individual users, can subsequently be used to authorize access to SharePoint resources.

When you create a Web application, you define the authentication method to use for identifying users who want to access content in the site collections that are within that Web application. A Web application can have only one authentication method associated with it. Therefore, if you want to access the same content via a different mechanism, you need to extend the existing Web application to another zone. There are five available zones: Default, Internet, Intranet, Extranet, and Custom. Each one is essentially just another IIS Web site configured with a suitable authentication method. The Default zone is automatically used when you first create a Web application, and the authentication method on this zone has to be Windows. The names of the other zones are simply descriptive and don't enforce any specific processing, but their names certainly reflect the most common scenarios in which you might deploy SharePoint content.

When you extend an existing Web application into another zone, you supply the URL for that zone. Thus different URL namespaces can invoke different authentication methods for the same content. We can see this in action in Figure 3, page 64, in which I'm attempting to access the same content via two different URLs—one using

Microsoft®

taking on dragons. easy.

1. Put the fire out.
Knowing what to do if there's a fire is always smart. That the fire spews from the mouth of a ferocious flying serpent should make no difference.

2. Ask for a break.
Searing heat, slashing claws, and the beating wings of hell will tire anyone. Say you need a break, then just walk quickly out the back.

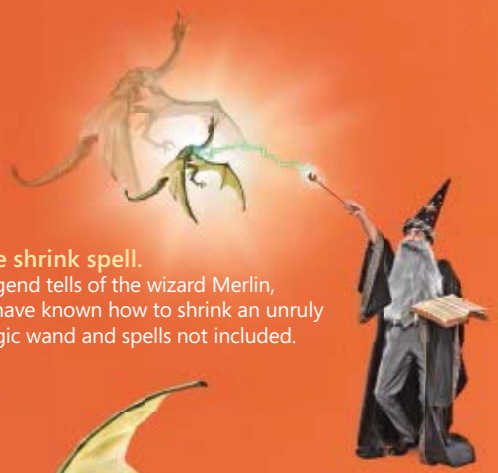

3. Use the shrink spell.
Arthurian legend tells of the wizard Merlin, who would have known how to shrink an unruly Dragon. Magic wand and spells not included.

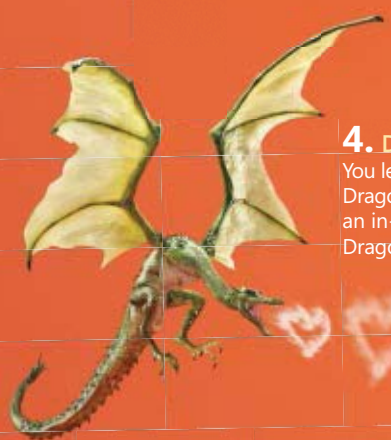

4. Dragonslayer.
You learn to slay Dragons by slaying Dragons. Win this one and you'll be an in-demand consultant to other Dragon-besieged companies.

5. The princess defense.
That temp in finance—bewigged, begowned, and pushed Dragonward—may just pass for a princess.

taking on security threats. easier.

1. Implement Microsoft® Forefront.™
Forefront makes defending your systems easier. It's a comprehensive, simple-to-use, integrated family of products that helps provide protection across your client, server, and network edge. Learn how Del Monte Foods uses the Forefront family of products to help defend their systems. Visit **easyeasier.com**

Forefront is business security software for client, server, and the network edge.

Microsoft®
Forefront™

Figure 3: Access the same SharePoint content via 2 different URLs

Windows authentication (http://employees.laphroaig.com) and one using forms with a SQL membership provider (http://friends.laphroaig.com). You can also use Alternate Access Mappings to support other namespaces, but that's outside the scope of this article.

With an extended zone with an authentication provider applied to it, you essentially have two sets of security principals that can now be used to authorize access to resources. When you define a membership provider, you give it a name, which is subsequently used to identify members within that provider. In this case, I have named my provider "Friends." You can see in Figure 4 that when I browse for users with a name of "Kevin," users from both the current Windows AD domain (ISLAY\Kevin) and the membership provider (Friends:Kevin) are returned. Now that you can identify who people are, let's see how you can grant them access to SharePoint resources.

Authorization

Authorizing access is a case of granting permissions to securable objects. When doing so, you need to consider the following five components:

Individual permissions. These permissions grant the ability to perform specific actions. For example, the View Items permission gives the user the ability to view items in a list. The list of individual permissions that are available are farm-wide but can be controlled at the Web application level by a farm-level administrator.

Permission level. This component groups individual permissions together for easier management and assignment. WSS has five default permission levels: Limited Access, Read, Contribute, Design, and Full Control. MOSS adds a few more such as Approve and Manage Hierarchy. You can add new permission levels or change the default levels to suit your needs. Permission levels are per-site and can either be inherited from a parent site or explicitly set at a subsite, library, or item level.

User. User is a security principal that can be identified using one of the authentication methods associated with the Web application.

Group. A group identifies a set of users. It can be a Windows security group, a role that's verified via a role provider, or a SharePoint Group such as Site Owners, Site Members, or Site Visitors. SharePoint Groups are new to SharePoint 2007 and essentially replace site groups. Groups provide a way for SharePoint site collection administrators to group users without having to rely on IT to create Windows security groups.

Securable object. Users or groups (either Windows Security groups, Roles or SharePoint Groups) are assigned a permission level for a specific securable object: site, list, library, folder, document, or item. By default, permissions for a list, library, folder, document, or item are inherited from the parent site, parent list, or parent library. However, anyone assigned a permission level that includes the Manage Permissions permission for a particular securable object can change the permissions for that securable object. SharePoint allows item-level permissions, which means that a user could be granted access to an individual document in a document library and not be able to access any other part of the SharePoint site. Access permissions on individual items also comes into play for the security-trimmed UI that SharePoint 2007 employs. You can see only items (including Web Parts) to which you have read access, and you can't see options whose function requires a permission that you don't have. For ease of maintenance, always use a group to assign permission levels to a securable object. Granting individual user access should be done only on an exception basis.

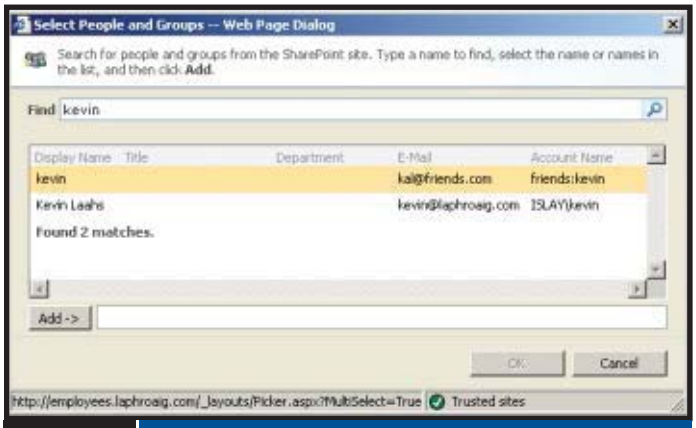

Figure 4: Using a membership provider called "friends" to search for members within that provider

Storing User Details and Establishing Permissions

So how does SharePoint store information about users such that it can subsequently vali-

date their access to resources? First, users can receive explicit access to objects via their user accounts or implicit access by being members of a security group or role. Furthermore, you can add users, security groups, or roles as individual entities or as members of a SharePoint Group. The latter is the preferred method as it eases overall management.

When you grant a user, security group or role any form of access to any resource in a site collection—either individually or via a SharePoint Group—an entry for that security principal is created in the UserInfo table in the content database that is associated with the site collection. (Their details are also put into the User Information List, which is what you see when you view All People through the browser interface.) Thus, if a security principal has access to multiple site collections, it will have multiple entries in the UserInfo table. This table stores, amongst other things, an internal identifier for the security principal that's used in many other tables, an indication of whether the principal is a security group or role, and the security identifier of the principal from its authentication provider. For the Windows provider, this is the Security Identifier (SID) of the user or group. For other providers, it's the unique identifier that that provider has allocated to the principal. If a user has been granted implicit access to a resource via a security group or role, then an entry in the UserInfo table is created for that user account the first time the user successfully accesses a resource in the site collection.

Four other tables come into play when establishing the permissions that a user ultimately has. First, the Groups and GroupMembership tables. When a SharePoint Group is defined, its details are stored in the Groups table, and the GroupMembership table has links to the individual users as defined in the UserInfo table. Thus, when you add security principals to a SharePoint Group, the GroupMembership table for that group is updated to include the internal identifiers for each principal in the UserInfo table. The other two tables are Roles and RoleAssigment. These are the tables that ultimately reveal the exact permissions that a requesting user has, with entries relating back to individual user records in the UserInfo table and SharePoint Group records in the Groups table.

Permissions that are associated with your individual user account, security groups, and roles of which you are a member and SharePoint Groups you belong to are aggregated to form your final list of permissions. Well very nearly; there's also something called Web Application Policies that apply permissions to the Web application as a whole, but detail for this is outside the scope of this article. Just know that these policies take precedence over permissions, and you can use them to globally deny or allow access to the Web application, then the individual user permissions come into play.

Concluding the Authentication and Authorization Process

So we now know how, through authentication providers, we can prove the identity of a requesting user and how we can store information about and assign permissions to known identities. So how does SharePoint validate that a requesting user is allowed to access a securable object?

If you pass the authentication process, SharePoint is passed the SID that identifies you in the underlying authentication scheme.

When you access any resource in a site you're essentially requesting an item from that site. For example, accessing the site's home page is actually making a request for default.aspx, and editing an item in a list is a request for EditForm.aspx. If you pass the authentication process, then SharePoint is passed the primary SID that identifies you in the underlying authentication scheme (e.g., your SID for the Windows scheme). This identifier is subsequently used to look up your details in the UserInfo table every time you request a resource. From there, SharePoint can establish whether you have the required permission to perform the task at hand.

It's important to note that SharePoint checks only your primary SID, which is important for the Windows security provider to know in cases in which your primary SID may have changed. This is typical in any form of domain migration.

Although many domain/user migration tools will retain the old SID in the user's security token, SharePoint doesn't check the sidHistory of the requesting user. Thus, there's no match in the UserInfo table for the new SID, and the user loses access. You can use the –migrateuser switch in the Stsadm utility to replace the old user account with the new user account, but you must take this behavior into account in your rename and migration processes to retain seamless access going forward. You can learn more about the Stsadm tool in "Stsadm: Taking Control of SharePoint Administration" November 2007, InstantDoc ID 97107.

Final Advice

SharePoint 2007 offers much flexibility in terms of authentication and authorization, which, when carefully planned, can result in a very robust and functional environment for sharing resources on many different levels and differing environments. Although you can use multiple authentication schemes, do be aware that there are some functional differences that can result. You can read about these differences in the article "Forms Authentication in SharePoint Products and Technologies (Part 3): Forms Authentication vs. Windows Authentication" (msdn2.microsoft.com/en-us/library/bb977430.aspx), and because of the way access checks are performed, be aware of the steps you need to take when migrating user accounts, so that users maintain access to their SharePoint content.

InstantDoc ID 98470

Kevin Laahs

(kevin.laahs@hp.com) is a principal consultant in the HP Services Advanced Technology Group. He is coauthor of *Microsoft SharePoint Technologies: Planning, Design, and Implementation* (Digital Press).

Q: What's the Active Directory Load Balancer tool?

A: With Windows 2000 Server, there's one bridgehead server per naming context, so under Win2K, no matter how many domain controllers (DCs) are in a hub site, only one would be used for intersite replication. This changed with Windows Server 2003, and intersite replication is now load-balanced over the available DCs in a location; however, this load balancing is performed only initially. If a new DC is added at a location, then existing connections won't be re-evaluated and spread more evenly between the DCs. Active Directory Load Balancer, which is part of the Windows 2003 Resource Kit tools, solves this problem by load balancing connection objects between the available DCs. The tool runs these three process phases:

- Gathers data about the sites and connection objects
- Calculates the new replication connection object design
- Writes the updated connection object properties

Load-balancing the connection objects between the DCs is useful, but the tool also staggers when replication will occur within the replication interval. For example, if it replicates site A between minutes 0 and 15 of the replication cycle, it then replicates site B between minutes 16 and 30. This is a great feature that spreads out the load experienced by the hub DCs, allowing them to support even more branch sites. Without Active Directory Load Balancer, all replication takes place in the first 15 minutes of a replication cycle.

InstantDoc ID 98415

—John Savill

At a Glance

In Windows Server 2003, if you have Authenticated Users, Creator Owner, Administrators, and System rights defined for C:\, can the Everyone group be removed without consequence from the default server build?

Yes, all user accounts on the system, except Guests, will fall into one or more of the groups Authenticated Users, Creator Owner, Administrators, and System. The Everyone group by default also includes everyone but Guests. The Everyone group is there to make it more convenient to set permissions. You should be able to remove it without any problems, although there's no real benefit in doing so.

InstantDoc ID 97945

—Mark Burnett

Q: How do I configure a domain controller (DC) to register site-specific records for an additional domain?

A: By default, a DC registers site-specific records for its own site. If you want a DC to also register records for an additional location (e.g., perhaps a location that has no DC of its own, and you want to control where the clients authenticate against), you can instruct the DC to register for additional sites. To do so, open the Group Policy Object Editor, go to Computer Configuration, Administrative Templates, System, Net Logon, DC Locator DNS Records, and use the *Sites Covered by the domain controller locator DNS SRV Records* Net Logon service Group Policy settings to specify the space-delimited site names for which the DC should register records, as Figure 1 shows. If you use Group Policy, you need to ensure that the GPO applies only to the DCs you want to register; therefore, you should apply a security filter to the GPO so that only specific DCs read the policy.

You can also add site names via the registry by updating the SiteCoverage value under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters registry key. Enter each site on its own line.

InstantDoc ID 98413

—John Savill

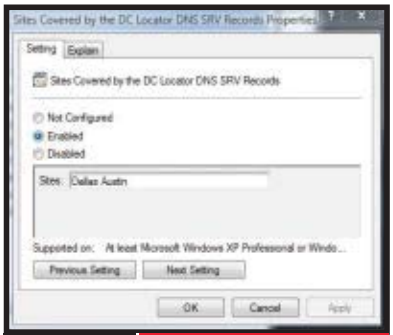

Figure 1: Specifying site names for which a DC should register records

Mark Burnett
(mburnett@xato.net)

John Savill
(jsavill@windowsitpro.com)

Control Windows Server 2008 Roles and Features

Eschew the Server Manager GUI and reach for the command line

Mark Minasi (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books, including *Administering Windows Vista Security: The Big Surprises* (Sybex). He writes and speaks around the world about Windows networking.

By the time you read this column, Microsoft will have released Windows Server 2008. Even if you've decided not to immediately upgrade to the new server OS, you're probably setting up a few test servers to try it out. And as you tinker with these fresh Server 2008 systems, you'll begin making observations.

First, Computer Management is gone! Right-clicking Computer and choosing Manage brings up not the familiar Computer Management console but rather the new Server Manager console. And if you try to use Server Manager to add, say, the DNS Server service, you'll see that you're now dealing with *roles* and *features*. Second, you can no longer use Control Panel's Add/Remove Windows Components applet. It doesn't exist in the new OS, so you'll be forced to use the newfangled Server Manager tool.

Pshaw, I say. Eventually, you're going to be faced with the prospect of rolling out dozens, hundreds, or thousands of servers and—let's all say it together—*you're not going to use the Server Manager console to do it*. Instead, let me introduce you to the tool's less celebrated command-line brother, ServerManagerCmd (servermanagercmd.exe), and its cousins Ocsetup (ocsetup.exe) and Oclist (oclist.exe). ServerManagerCmd works only if you're running the "full" (GUI-equipped) version of Server 2008. (I'll get to the Server Core commands in a bit.)

ServerManagerCmd

Working from an elevated command prompt (as always) in Server 2008, type

```
servermanagercmd –query
```

This simplest of ServerManagerCmd commands will display several screens that show all the add-on capabilities that your server currently has and hasn't enabled. (I say "capabilities" because Server 2008 breaks them into roles and features, and I haven't figured out the difference. Given that DNS is a role and WINS is a feature, can we assume that roles are cool and features are less so?) One such line might look like

```
[ ] DNS Server [DNS]
```

where *DNS Server* means that the server has the potential to become a DNS server, the empty opening brackets mean that you haven't enabled the service (if you did enable it, an X would appear between the brackets), and—most important—*[DNS]* informs ServerManagerCmd that you want to perform a task involving the DNS server service.

That terminology is the key to making this server capable of acting as a DNS server:

```
servermanagercmd –install dns
```

Notice that ServerManagerCmd recognizes the DNS role but is case-insensitive.

Ocsetup and Oclist

Now, if you're setting up DNS on a Server Core installation of Server 2008 rather than a full installation, you'll find no ServerManagerCmd tool; instead, to list the capabilities you do and don't have installed, you can type

```
oclist
```

You'll see output similar to that which ServerManagerCmd -query provides, but the official names for the roles and features are lengthier. For example, the DNS line looks like

```
Not Installed:DNS-Server-Core-Role
```

To install a role or feature, just use Ocsetup, followed by the role or feature's name. To install DNS, for example, you'd type

```
ocsetup DNS-Server-Core-Role
```

Note that you don't need an -install option. Ocsetup only installs and uninstalls; it has no option like -query because the separate Oclist program fills that need.

You'll also find that, like many Server Core–specific commands, Ocsetup is one of those "strong, silent type" commands. If you've typed an Ocsetup command with correct syntax, your only response is another command prompt. Be aware, though, that just because you've gotten another command prompt doesn't mean that Server Core has finished installing the DNS server role. Occasionally, I've gotten in trouble by assuming that. Now, I use the belt-and-suspenders technique of running Oclist on a Server Core system to be certain that my desired role has been installed.

Command-Line Power

Honestly, would you rather click through Server Manager to install DNS when you can just type

```
servermanagercmd –install dns
```

from the command line? Of course not! So go out and buy one of those comfortable keyboards, because Server 2008 is here!

InstantDoc ID 98303

Vista Command-Line Tools

Use these tools to mirror directories, manage ACLs, perform backups, and complete other important administrative functions

Michael Otey (mikeo@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and coauthor of *SQL Server 2005 Developer's Guide* (Osborne/McGraw-Hill).

The lion's share of the attention to Windows Vista always seems to center around the new interface, which is understandable considering its exciting new look and feel. Unfortunately, part of the attention is because you now have to find new ways of doing old tasks. However, new features in Vista don't end with the UI. Under the covers, Vista also sports command-line tools—many of them new—that make a powerful addition to your administrative toolbox. Here are my favorites.

10 Bcdedit—The days of editing the simple boot.ini file are gone. Vista's new boot process saves its system boot configuration in the Boot Configuration Data (BCD) store. Like all bad ideas, the BCD store replaces a simple concept with a complex albeit more secure one. Bcdedit is your primary tool for editing the BCD store. Bcdedit supports a wide set of command-line options. For instance, to list the contents of the store you can run

```
bcdedit /enum
```

9 Choice—A handy batch-file command, choice lets you display a list of choices to users of a command-shell script. The choice command returns an index value in the ERRORLEVEL environment variable indicating the user's selection. For example, the following command prompts users to enter Y, N, or C; the ERRORLEVEL variable returns 1, 2, or 3 respectively:

```
CHOICE /C YNC /M "Press Y, N, or C."
```

8 Waitfor—The waitfor command is a useful scripting command that synchronizes processes running on multiple systems on the network. As its name suggests, waitfor can pause a script until the command processor receives a specified signal, and it can also send a signal to one or more systems on the network. The following command waits for the ScriptDone signal:

```
waitfor ScriptDone
```

7 Wbadmin—Although you might not have a clue based on its name, the new wbadmin command is Vista's command-line backup tool. The following example shows how to use wbadmin to backup the C and D drives to the share named backup on myserver:

```
wbadmin start backup -backupTarget:\\myserver\
  backup include:c:,d:
```

6 Icacls—The Icacls command replaces the older cacls command. Icacls lets you list, update, and back up the access control lists (ACLs) for files and directories. The following example shows how you can save the ACLs for the C:\temp directory:

```
icacls c:\temp /save tempacl
```

5 Winsat—Winsat is the new Windows System Assessment tool. It runs automatically when you install Vista, but you can run it on demand for simple system benchmarking and system information. For instance, to list your system information, you could run winsat with the features parameter:

```
winsat features
```

4 Clip—The new clip command is a handy tool that copies the output of other command-line tools to the clipboard. The following example shows how to use clip to copy the contents of the file mytext.txt to the clipboard:

```
clip < mytext.txt.
```

3 Forfiles—The forfiles command is another useful batch-file tool. It executes a command over a set of files. Forfiles is much easier to use than the older and more obtuse for command. The following example shows how you could list all files more than 30 days old in the c:\temp directory:

```
forfiles /p c:\temp /s /d -30 /c "cmd /c echo @file
```

2 WinRS—The WinRS tool is essentially the Windows version of Linux's Secure Shell (SSH). WinRS lets you open a secure command window to a remote host. All the contents of the remote shell are encrypted. The following example connects to the server named myserver and runs the dir command:

```
winrs -r:myserver dir
```

1 1. Robocopy—Without a doubt, the best command-line addition to Vista is robocopy. Although robocopy isn't new to most of our readers (it's been a staple in the Windows Resource Kit for years), Vista is the first release that includes robocopy as a part of the OS. Robocopy is a super-powerful command-line copy tool. The following example shows how to create a mirrored copy of the directory called shares and all of its subfolders.

```
robocopy "C:\Shares" "\\server2\ Shares Backup" /MIR /R:2 /NP
```

InstantDoc ID 98458

Did You Know...

Along with windowsitpro.com and sqlmag.com two new sites have been launched to ensure custom-made content is just a click away.

Office & SharePoint PRO .com
Formerly MSD2D.com

Microsoft Office and SharePoint content mentored by a community of peers and professionals.
www.officesharepointpro.com

windowsdevpro.com

A community addressing the need of content for the developer who needs to create with the IT administrator in mind.
www.windowsdevpro.com

Windows IT Pro

Engage with our network of peers and professionals and view various forms of content. It is a complete source for IT Professionals and managers.
www.windowsitpro.com

SQL Server magazine

For information on managing, mining, building and developing world-class applications.
www.sqlmag.com

Jeff James (jjames@windowsitpro.com)
is senior editor, products, for *Windows IT Pro* and *SQL Server Magazine*.

Readers Review HOT PRODUCTS

At a Glance

File Copy Utility

Microsoft SyncToy 2.0 (Beta)

—Jeff Kowalke, senior Microsoft systems engineer

Reader: Jeff Kowalke Senior Microsoft Systems Engineer
Product: Microsoft SyncToy 2.0 (Beta)
Company: Microsoft
Contact: www.microsoft.com

As many IT pros know, some of the best tools and utilities to help make their jobs easier are freeware and shareware products. Such is the case with Senior Microsoft Systems Engineer Jeff Kowalke, who was often asked by end users for a quick backup solution. Kowalke had been relying on the robocopy command in script files, but wanted something easier to use.

After some research, Kowalke came across the beta version of **Microsoft SyncToy 2.0**, a file copying and synchronization tool. "I recently upgraded my home computer, and I had to search for old digital photos scattered across a wide variety of sources, like my old computer, old DVD backups, and backups saved onto an external USB disk," says Kowalke. "SyncToy helped by finding sub-folders of those backups that I would have ignored…SyncToy has been a true time saver."

The ability of SyncToy's backup feature to keep files in their original format after backup operations is one of Kowalke's favorites, a feature that avoids end users dealing with .BKF or .ZIP file compression formats that they may not have the correct software to uncompress.

"Since SyncToy is considered a free and unsupported utility from Microsoft, I really can't complain about it," says Kowalke. "I'm happy that Microsoft is willing to make this tool available for free to fill an obvious need for a simple, visual backup utility. That said, I would [suggest that Microsoft] add the ability to schedule backups, and add a reporting feature that would display backup results via email when a backup completes. I could then use this product in a server environment for disk-to-disk backups."

What's Hot continues on page 72

Wanted: Your Real-World Experiences with Products

Have you discovered a great product that saves you time and money? Do you use something you wouldn't wish on anyone? Tell the world in a review right here in What's Hot: Readers Review Hot Products. If we publish your opinion, we'll send you a Best Buy gift card and a free online subscription to a ProVIP publication of your choice! Send information about a product you use and whether it helps you or hinders you to whatshot@windowsitpro.com.

.NET Performance Optimization

Strangeloop Appscaler

Reader:
Jackie Reid
VP of Client Services

Product:
Strangeloop AS1000 Application Scaling Appliance

Company:
Strangeloop Networks

Contact:
www.strangeloop.com

Improving the performance of ASP.NET and AJAX applications is a task that many web developers often struggle with, just as web marketing technology developer Marqui has. In order to improve the performance of their .NET applications and hosted Web sites, Marqui decided to give the **Strangeloop AS1000 Application Scaling Appliance** a try.

According to Marqui VP of client services Jackie Reid, their primary goal was to improve the customer experience by improving the speed of their web applications. "We're using two AS1000 products, and we're currently treating web traffic on our Marqui 5.0 Application Server and a main hosting server," says Reid. "When the [next AS1000 OS upgrade] becomes available that offers failover redundancy and analytics, we'll be moving the AS1000s in front of all of our Web traffic."

Reid says that deployment was straightforward, and that they were up and running on the appliances within roughly an hour. After deploying the AS1000s, Reid says that they've experienced "significant performance improvements" in both developed ASP.net applications and on Web sites they host for clients.

Colin Edwards, Marqui's manager of client operations, likes the AS1000's ability to create custom data filters, an ability that lets him "really control what's going on in my data center." As useful as those filter treatments are, Edwards does have some advice for anyone that may use them. "One of the treatments I created caused a problem with the traffic going out of the data center," says Edwards. "I talked to [Strangeloop] support to find out what I did wrong [and to resolve the problem]. You should really know what you're doing before experimenting with different filter treatments."

What's Hot continues on page 76

VM PERFORMANCE MONITORING

Just Point, Click & Know How Your Virtual Machines Are Performing

When running hundreds to thousands of virtual machines and multiple VirtualCenters, it is imperative to deploy a centralized and scalable monitoring solution to quickly pinpoint trouble spots and easily track the impact of higher hardware utilization and virtual machine performance.

vCharter Pro™ not only helps users prevent unwanted downtime, it can also assist in optimizing resource allocation among virtual machines – a major consideration for companies deploying large-scale virtualization implementations.

For more information about vCharter Pro™ or our full lineup of software, visit us at www.vizioncore.com.

08198

Microsoft altiris LANDesk

KACE

Deploy in days, not months.

No kidding around. Installing a KBOX by KACE gives you complete systems management in days, not months. We'll also do it for one-third the cost of the big three. Give us a call, let us prove it. Welcome to KACE Time.

EMA ALL-STAR 2008

Enterprise Management Associates
2008 Rising Star

KACE™
Systems Management. Done.

www.kace.com/deployindays 877.MGMT.DONE

Systems Management

NetIQ Aegis

Reader: Neil Forster Network Engineer
Product: NetIQ Aegis
Company: NetIQ
Contact: www.netiq.com

Outsourcing IT and Web operations to a third party has become a viable option for many companies, a development that plays into the strengths of Attenda, a UK-based company that specializes in providing IT management services. In the interest of improving the services that they offer to customers, Attenda recently adopted the **NetIQ Aegis** IT automation platform.

"As a company we've been a long-term user of NetIQ's AppManager product," says Neil Forster, a network engineer at Attenda. "We've also been looking at the run book automation (RBA) market for some time, and we knew that NetIQ would be coming out with something in the RBA market, which eventually turned out to be Aegis."

According to Forster, Attenda decided to implement Aegis in order to make their IT management services more efficient. "One of our concepts is operational certainty, [which we define] as the ability to do something the correct way over and over again. [NetIQ Aegis] gives us the opportunity to automate responses that have historically been handled manually…such as taking help desk calls from notification to resolution."

Forster likes the ability to easily create basic process workflows, and believes that adoption of the NetIQ Aegis system has helped them make some of their existing processes more efficient. "For example, we used to have a manual procedure that would require an engineer to look for particular events in an event log, check a registry key value, and then reset that value based on the log information… we now have automated that process using NetIQ Aegis."

Basic process workflows may be easy to configure and use, although Forster points out that creating more complex process automation tasks may take some additional time. "Configuration of more complex processes wasn't entirely intuitive," says Forster.

"We used to have a manual procedure that would require an engineer to look for particular events in an event log, check a registry key value, and then reset that value based on the log information. We now have automated that process using NetIQ Aegis."

—Neil Forster, network engineer

"We also had some problems with performance, particularly with an environment with a large number of servers. We did work closely with NetIQ to address those problems, and they provided hot fixes for those."

InstantDoc ID 98584

WANTED

FOR DATA BREACHES, IDENTITY THEFT, AND HARBORING SPYWARE AND VIRUSES.

CIOs and IT personnel are at risk of losing vital information and data, and are advised to search for alternative computing methods. Desktop PCs and laptops are prone to data breaches, hackers, spyware, viruses, and other crippling problems that can destroy IT infrastructures everywhere.

Devon IT's line of thin client terminals make data theft virtually impossible. Data is stored and managed on your enterprise servers and can only be accessed by authorized users. Thin clients provide true PC experience without the threats of data theft and robbery.

Visit www.devonit.com/wanted or call 1.888.524.9382 for more information, or email info@devonit.com to receive FREE White Papers and Case Studies about how thin clients have helped protect companies across the world.

SafeBook Notebook – Where Security Meets Mobility

- No hard drive, so no sensitive data can be lost
- Runs anywhere, through wireless, Ethernet, or 3G Broadband connections
- Battery lasts for over 6 hours
- HIPAA Compliant
- **Starting at $599**

High Powered Desktop Access Device

- Fastest thin client in the world
- Increased security
- Low total cost of ownership
- Fanless

DEVON IT

Server-Based Computing for the Modern Business

www.devonit.com

DIRECTORY OF SERVICES

Windows IT Pro Network

Search our network of sites dedicated to hands-on technical information for IT professionals.

www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.

www.windowsitpro.com/forums

News

Check out the current news and information about Microsoft Windows technologies.

www.wininformant.com

EMAIL NEWSLETTERS

Get free NT/2000/XP/2003 news, commentary, and tips delivered automatically to your desktop.

- *Windows IT Pro UPDATE*
- *Vista UPDATE*
- *Windows Tips & Tricks UPDATE*
- *WinInfo Daily UPDATE*
- *.NET Briefing*
- *Exchange & Outlook UPDATE*
- *Scripting Central*
- *Security UPDATE*
- *SQL Server 2005 Express UPDATE*
- *SQL Server Magazine UPDATE*
- *Windows IT Library UPDATE*
- *Connected Home EXPRESS*

www.windowsitpro.com/email

PRO VIP ACCESS

Exchange & Outlook Pro VIP

Discover smart solutions for Exchange and Outlook administrators.

www.exchangeprovip.com

Scripting Pro VIP

Learn how to create more powerful scripts and get tips for automating those tedious administrative tasks.

www.scriptingprovip.com

Security Pro VIP

Discover practical, how-to advice for avoiding and solving security problems.

www.securityprovip.com

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Contact Joel Kirk at jkirk@penton.com.

Super CD/VIP

Get exclusive access to all of our print publications, including *Windows IT Pro,* via the new, banner-free VIP Web site.

www.windowsitpro.com/sub/vip

Article Archive CD

Access every article ever printed in *Windows IT Pro* magazine since September 1995 with this portable and speedy tool.

www.windowsitpro.com/sub/cd

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.

www.sqlmag.com

For detailed information about products in this issue of *Windows IT Pro*, visit the Web sites listed below.

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

SEND US YOUR INDUSTRY HUMOR! Email your funny screenshots, favorite end-user moments, and humorous IT-related pics to rumors@windowsitpro.com. If we use your submission, you'll receive a Ctrl+Alt+Del coffee mug.

RAD commentary!

Totally, Dude

We prefer you just leave it in the package

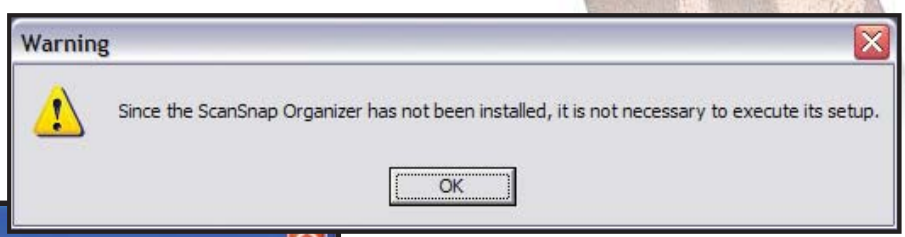

IF ONLY it were the fate of conference rooms

The IT Pro at Home!

What do you do after a long day at the office tinkering with systems and dealing with end-users? We're willing to bet you go home and do the same thing! You tinker with your home-networking setup, share media files across your systems, and solve the problems your family members are having with their satellite systems. You've got a connected home, and you probably use many of the same solutions there as you do at work. That's where *Connected Home Media* (www.connectedhomemag.com) can help. You're not only the IT pro at work—you're the IT pro at home! Sign up for the free *Connected Home Express* newsletter (www.windowsitpro.com/email) and get your tips about media sharing, home-network security, backup and recovery, home theater, and more!

May 2008 issue no. 165, *Windows IT Pro* (ISSN 1552-3136) is published monthly. Subscriptions in US, $54.95 for one year; in Canada, $59 US currency, plus GST for one year; in all other countries, US $99. *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538, (800) 793-5697 or (970) 203-2782. Sales and Marketing Offices: 221 E. 29th St., Loveland, CO 80538. Advertising rates furnished upon request. Periodicals Class postage paid at Loveland, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80539-0447. SUBSCRIBERS:** Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 221 E. 29th St., Loveland, CO 80539. Printed in the USA. BPA Worldwide Member.

SERVERS, STORAGE, SWITCHES. ONE BOX. NO HASSLES.

intel Xeon inside™ Powerful. Efficient.

IBM BladeCenter S Express

$4,499 (SAVE $493)

OR $119/MONTH FOR 36 MONTHS[1]

Introducing IBM BladeCenter S Express. Now you can combine Quad-Core Intel® Xeon® processor–based blade servers, storage, switches and management tools in one small chassis. It's easy to set up. Easy to use. Easy to manage. It's a simple way to simplify your IT.

From the people and Business Partners of IBM:

It's innovation made easy.

SIMPLIFY AND MANAGE YOUR I.T. WITH A SINGLE CHASSIS.

PN: 8886E1U

- Up to six application blades with the ability to expand to multiple virtual blades
- Integrated storage built into the chassis – 3.6TB SAS or 6TB SATA
- 3-year customer replaceable unit and on-site limited warranty[2]

IBM BLADECENTER HS21 EXPRESS

$2,359 (SAVE $249)

OR $62/MONTH FOR 36 MONTHS[1]

PN: 8853E1U

- Features up to two high-performance Dual-Core or Quad-Core Intel Xeon Processors
- 1GB standard/16GB maximum memory per blade (32GB with Memory and I/O Expansion Unit)
- 3-year customer replaceable unit and on-site limited warranty[2]

IBM SYSTEM STORAGE DS3300 EXPRESS

$4,545 (SAVE $450)

OR $120/MONTH FOR 36 MONTHS[1]

PN: 172631E

- Support for dual-port and hot-swappable SAS disks at 10,000 and 15,000 RPM speeds
- Expandable by attaching up to three EXP3000s or a total of 48 hard disk drives
- 3-year limited warranty on parts and labor[2]

IBM Express "Bundle and Save"

We bundle our Express systems to give you the accessories you need – while saving you money on the hardware you want. Act now. Available now through ibm.com and IBM Business Partners.

IBM express advantage™

ibm.com/systems/onebox

1 866-872-3902 (mention 6N8AH01A)

1. IBM Global Financing offerings are provided through IBM Credit LLC in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government customers. Monthly payments provided are for planning purposes only and may vary based on your credit and other factors. Lease offer provided is based on an FMV lease of 36 monthly payments. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.
2. IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. For a copy of applicable product warranties, visit ibm.com/servers/support/machine_warranties or write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services, including those designated as ServerProven® or ClusterProven.® Telephone support may be subject to additional charges. For on-site labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician. On-site warranty is available only for selected components. Optional same-day service response is available on select systems at an additional charge. IBM, the IBM logo, IBM Express Advantage, IBM BladeCenter, System x and System Storage are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM trademarks, see ibm.com/legal/copytrade.shtml. Intel and Xeon are registered trademarks of Intel Corporation. All other products may be trademarks or registered trademarks of their respective companies. All prices and savings estimates are based upon IBM's estimated retail selling prices as of August 1, 2007. Prices and actual savings may vary according to configuration. Resellers set their own prices, so reseller prices and actual savings to end users may vary. Products are subject to availability. This document was developed for offerings in the United States. IBM may not offer the products, features, or services discussed in this document in other countries. Prices are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or IBM Business Partner for the most current pricing in your geographic area.

FEEL LIKE YOU'RE STORING EVERYTHING AND MANAGING NOTHING? BROCADE FILES MANAGEMENT SOLUTIONS HELP YOU TAKE BACK CONTROL.
With Brocade Files Management Solutions, you can automatically and transparently migrate files to the optimum types of media based on your rules. So you can dramatically lower data management costs and gain more control of your file environment without compromising users' needs. And get a lot more breathing room. Get your free eBook on File Data Management at: www.brocade.com/take_control

BROCADE